Remote

After years of hand wringing and false starts, remote work finally has its moment. Suddenly, massive companies are being run from breakfast nooks and messy bedrooms. Those first Zoom [1] meetings were pretty rickety–half the employees didn't know how to log in, and the other half didn't think they needed to wear shirts. It sometimes seemed that a huge percentage of the internet's bandwidth must be carrying nothing but the sounds of crying children and the deafening feedback of people who cannot seem to learn that they can't have two microphones on at the same time[2].

This episode explores the research and realities of a fully remote world. In "Business-as-Usual" we review the limited existing research on remote work and find that many of our assumptions about distractions and productivity don't hold true. "Measuring Remote Success" explores new metrics that take us from in-office productivity to employee growth. The research on remote work reveals how different personalities require different support; "People Are Different" identifies core remote worker profiles and what they need for success. But it's not all about individual employees–"Controlling the Beyond-Control" begins exploring the vital role that companies have in supplying what their employees cannot. "Rebuilding Culture" reveals the cultural and organizational factors that predict success for distributed teams, and identifies the technologies native to distributed work. Finally, "Distributed Innovation" and "Diversity-Innovation Paradox" confront two of the biggest mysteries of remote work: innovation and inclusion. Without change, neither are possible, but these sections present a new framework that marries innovation and inclusion to move beyond the lazy limitations that have held us back.

Business-as-Usual

Outside a handful of quirky mid-sized tech companies, almost no one has run a completely distributed workforce before. It is quickly becoming clear that we have no idea what we are doing. One large-scale Harvard survey found that only 40% of companies felt they were "well prepared" for "flexible work", much less a business-wide shift to remote work. Even the US Supreme Court has had to adapt to the new remote (it seems that not all of the "attendees" understood how to use mute). Unfortunately, there is very little research on universal remote work.

---

[1] Zoom: a verb–to use *Meet*, *Teams*, *WebEx*, *GoToMeeting*, or other video conferencing service for a meeting. Coincidentally, there is a company called *Zoom*. Lucky them. (It certainly can't be that the entire world is made up of people that can't tell the difference between a product and the very, very not new idea of video conferencing. That would be like when my mom says that the "internet" isn't working, and then I must spend the next hour figuring out if she's talking about her browser, router, email, computer, SMS, phone, Facebook, wi-fi, or some other unrelated concept that in her mind is part of the singular entity called "internet".

[2] Are these the same people that routinely set off their own car alarms when heading off to work at 6AM?

While we may not know much about what happens when an entire company shifts to working from home, there has been extensive research done in two domains which might help us: distributed computing and distributed cognition. Both deal with the problem of how to do something (search a database or make a judgment) across a collection of distinct nodes (computers or brains). Combining these domains not only helps us understand the flaws of remote work but points to something entirely new: distributed work.

The modern world could not function without distributed computing, but distributed computing is more than just a bunch of PCs operating in different locations. It is an entirely different approach to computing that is adapted to the "fragile narrow laggy asynchronous mismatched untrusted pipes" connecting those CPUs together. Similarly, distributed work requires much more than a bunch of individual, pajama-clad employees working from their bathroo…bedrooms.

Distributed work requires different behaviors to be successful; using remote work as a patch without any adaptive changes doesn't automatically translate for every company or employee. For example, the CEO of one fully distributed company, Automattic, observes that remote workers at most companies still expect "hyper-responsiveness" from their coworkers, who must continue to maintain regular hours and other maladapted habits inherited from traditional office work. He claims these bad practices "suppress the ability of knowledge workers to actually think".

Research supports his claim. Hyper-responsiveness and multitasking across messaging and media tools reduces cognitive control and decreases cognitive performance by 60% in complex, creative work. Employees often work faster to compensate for increased interruptions but experience "more stress, higher frustration, time pressure and effort".

During the early months of "remote work" average productivity increases, but over time most workers lose those gains and even fall below pre-remote levels. Only a small group remains hyper-productive and engaged. This drop in productivity has been hidden in much existing research as working from home has been "a privilege that you earn" at most companies; therefore, the studied workforce was self-selected to be self-motivated.

Many consultants in this space have argued the benefits of fewer hours wasted in meetings and more time is spent on "getting things done"; however, there is surprisingly little research on remote work supporting those claims. While some studies find employees self-report fewer interruptions, there is no objective confirmation. Do onsite interruptions outweigh the increase in digital and nonwork interruptions? As one researcher noted, "In person…the social cost of asking someone to take on a task is amplified… In a remote workplace, in which co-workers are reduced to abstract e-mail addresses or Slack handles, it's easier for them to overload each other…" This dissonance in claims about remote work likely reflects the different types of work being observed (e.g., routine vs. creative). Organizations need objective measures of interruptions and their effect rather than guesswork and introspection.

Many business leaders have debated whether the current situation of Covid-19 lockdowns will accelerate the shift to digital and collaborative tools. It's possible, but this acceleration may only be true for companies already in the process of transition. In this episode, we will explore how everyone can benefit from a transition to distributed work, from diverse individuals to whole companies, from daily routine to the cutting edge of innovation.

Creativity isn't finding a solution to a given problem. It's what happens when you stop waiting for someone to give you a problem and start exploring the unknown. Offices never forced us into passive, hierarchical work cultures, but they allowed them to persist even as economic reality has changed. Instead of adapting remote technology to traditional work practices, we need to adapt our practices and technology towards a true distributed work enterprise.

Measuring Remote Success

Writing in the New Yorker, Computer Science professor Cal Newport drew an analogy between the transition from centralized to remote work and the transition of factories from centralized steam engines to distributed electric motors. He argues, "…most companies that have tried to graft it onto their existing setups have found only mixed success," in large part because mid-level management has been resistant to change. Electric motors demanded entirely new configurations of factories. Networked computing and distributed algorithms forced even more radical departures from the past. Abandoning the centralized workplace will require an even greater transformation, as if you're launching a competitor, but one already native to distributed work.

One of the principal challenges in distributed computing is that it is fragile: a failure at any point in the network can block the entire system. Computers can only respond to contingencies for which they have been explicitly programmed[3], and so every individual piece must be engineered to handle all possible failures. Because distributed computing might have many different components, identifying a point of failure in the moment can be nearly impossible and often requires sophisticated analytics to monitor everything.

By contrast, a point of failure on a colocated team (e.g. a sick day) is easily identifiable, and the entire rest of the team can quickly adapt with minimal preparation. A hierarchical team doesn't need to be programmed to handle many contingencies because leadership can quickly redirect work. When working remotely, such a failure is much slower to reveal itself. People aren't visibly absent, and their contributions will not be automatically absorbed by others. Leaders also cannot easily redirect the flow of work as people are simply not as responsive or available. This might suggest a need for deep intrusive work analytics to constantly monitor the productivity of

[3] This may not always be true as artificial intelligence doesn't need to be explicitly programmed to do anything, though if it is exposed to a contingency it has never seen before, spectacularly unpredictable things have been known to happen.

every member of a team, but people are not computers and they do much more than execute an algorithm.

While remote work[4] can be fragile, we know from distributed cognition that people can adapt in ways computer systems cannot. As people work together to solve problems they develop communication shorthands, novel efficiencies, and even learn to model one another's thought processes. All of this allows the individual nodes of distributed work (i.e. employees) to adapt to hidden failures throughout the network, sometimes without even being aware of it. This power of distributed cognition is a huge advantage over distributed computing, and companies should explicitly train for this adaptability. Don't treat human beings as dumb points of failure in a fragile system that need intrusive, ubiquitous monitoring. Use work analytics to understand how to increase the robustness and adaptability of your employees.

Traditional office work makes use of productivity metrics and KPIs such as net sales, profit margin, and order fulfillment time. The gig economy, which has invested heavily in remote work, has obscenely expanded these same metrics, logging individual keystrokes and time-on-task to the second. However, we already know that most employees begin to struggle with remote work at some point. Decreases in those numbers aren't a measure of an employee failure so much as a diagnosis of a process failure. The transition to distributed work requires KPIs that measure adaptability rather than raw productivity, and even better, analytics that provide a map for how to get there. (As we will discuss in the next section, that map will be different for different people.)

Some researchers have begun exploring these new analytics. One finding reveals that "conspicuous monitoring", transparently monitoring one specific task domain, "improves performance on task dimensions not being directly paid for." Rather than logging every keystroke and tracking every dimension of work, simply visibly attending to employees promotes broad increases in productivity.

New Metrics for Growth

In our analysis of social tracking technologies for Covid-19, we discussed how analyzing social graphs can allow us to predictively isolate potential superspreaders. My inspiration for this project was my previous research on how work-based social interactions between employees affected productivity. The question is less, "How many times were you interrupted in a day?" and more, "Which interruptions increased productivity and which were harmful?" While few offices are instrumented to measure such interactions, in distributed work this information is both readily available and crucial. A number of metrics can be starting points.

Remote interaction variables:

- Number
- Duration

[4] *Remote work*: forcing traditional work practices into remote environments. (We can do better.)

- Relationship
    - e.g., manager, co-worker, direct report, inter-team
- Context
    - e.g., meeting, planned, unplanned, one-on-one, group, time
- Medium
    - email, messaging, video

Post-interruption productivity variables:

- Individual level productivity
- Team level productivity

Possible models of the relationship between interruption and productivity:

- Simple correlation between productivity and types of interruption
- More advanced time-varying social graph clustering of factors causing changes in productivity[5]

Use these new (and old) variables to answer the following questions:

- Does distributed work actually come with more interruptions?
    - Electronic interruptions?
    - Non-work interruptions?
- How does this relate to employees' perceptions about interruptions?
- Which types of interruptions increase productivity?
    - Frequency patterns
    - Type of relationship
    - Context
    - Nature of work
        - creative vs. routine
- Which types of interruptions decrease productivity?

Beyond raw productivity and interruptions, there's also the problem of recognition. People that work remotely receive less recognition for their contributions and receive fewer promotions. Those telecommuters that do receive recognition have more face time with their managers. Remote work is a classic case of "out of sight, out of mind", and we know that promotions have always been a function of proximity. Distributed work analytics must actively balance engagement by managers to prevent promotion based on the *availability heuristic*.

Companies can't rely on existing productivity measures to define success in a distributed world. Develop new measures to help discover the paths to employee growth rather than snapshots of raw productivity. Don't just change what you are measuring, but why you are measuring.

---

[5] We'll return to social graph analysis and distributed cognition in "Distributed Innovation" and it will be wild… nerd wild!

Some of the above analytics may seem complex while others are simple adjustments. In either case, we must shift from prioritizing old metrics and towards metrics that understand employee success in a distributed workforce. Focus on why something isn't working for certain individuals rather than punishing people for decreased productivity. Balance near-term productivity needs against building longer-term human capital capacity and consider long-term aims. Is what you're doing now bettering your company post-lockdown?

People Are Different

In distributed computing, one of the challenges is that the components of the system might be wildly mismatched. Cutting-edge processors interact with outdated hardware, and the latest operating systems receive packets from ancient versions. Distributed computing must be engineered to handle all of the different hardware and protocol versions in its network. Rules must be explicitly established to allow these systems to work together.

In distributed cognition, those rules often naturally evolve as individuals interact, developing into norms, shorthand, and shared culture. In fact, distributed cognition is more than just a metaphor; the brains of high-performing teams of humans literally sync up. In classrooms, for example, students whose brains show less synchrony with other students perform worse, and brain imaging of engineers reveals shared neural patterns in response to engineering-related scenarios. Whether computing or cognition, mismatch can throw a distributed system out of alignment.

Obviously, it doesn't take remote work for people with differing capabilities and work styles to become frustrated working together. Working remotely, however, transforms some of these experiences. On the negative side, remote work can exacerbate the effects of differing cultural norms. On the upside, an unexpected benefit from early surveys in Los Angeles suggest that public employees working remotely engage in fewer downward comparisons. In other words, they spend less time complaining about their co-workers' laziness.

One of the few core findings we already have concerning remote work is that some people will be successful working remotely and some won't. Most studies of companies' remote work practices haven't revealed this because remote work has largely been a small, earned privilege. Now that many companies have gone entirely remote, differing populations will need different support in order to be successful and productive in the context of distributed work. We simply can't assume that standard onsite methods for supporting employees are still viable.

One challenge for understanding individual differences is the measurement itself. For example, more *neurotic* individuals (as measured on the Big Five personality inventory) report a preference for working from home but are in fact less successful. It is emotionally stable individuals that are resilient to the challenges of remote work. Similarly, *conscientious* individuals prefer the idea of working from home, possibly in response to the popular notion that

fewer distractions will allow them to get more done. Despite those expectations, people that have actual experience with the unique challenges of remote work generally haven't preferred it. While actual data on distractions is hard to come by, this preference against remote work by people who have done it suggests hidden obstacles.

So, the research indicates that we need to find synchrony between mismatched coworkers, but it also reveals that we understand much less about remote work than we assumed. If we wish to take full advantage of distributed work, we must start by understanding the mystery of why a minority of employees are robustly successful while most struggle, and develop targeted practices that support each group. This turns out to be a question of balance versus synergy[6].

Balancers vs. Synergists

The majority of employees rely on the structure, explicit and implicit, of a traditional work day. Much of how they work flows from the daily rhythms of arriving at an office, participating in regular meetings and interactions, and simply being present in the work environment. In fact, the majority of workers desire a clear distinction between their work life and their personal life with unambiguous boundaries. For them, the biggest challenge is establishing a work-life balance. These "balancers" are that majority group that has struggled with remote work.

For the rest, work is part of their core identity, and stripping away the barriers between home and office allows an even greater synergy between their work and life. This group has thrived working from home (as they often thrive in any work environment). The biggest threat to these "synergists" isn't a lack of structure but, ironically[7], too much.

Previous research has identified this split between balancers and synergists in traditional work environments, exploring topics such as boundary management and micro-transitions. The contrast between synergists and balancers isn't "good" versus "bad" employees. Traditional workplaces scaffold most employees' self-management, providing support for what is commonly a weakness. Yet even in traditional workplaces, our research finds that a broad array of cognitive, affective, and social skills–we call this *Meta-Learning*, learning how to learn–are more predictive of quality of work than education level or work skills. Now that the scaffolding of the workplace has suddenly been removed, it is meta-learning, not elite university degrees, that drives those hyper-productive synergists. Ignore intuitions that those with education or seniority will all be synergists, and hide your surprise when some lower-level employees find success in a distributed world.

---

[6] Yes, I know that lexical honesty of "synergy" has been greatly diminished by business newspeak, but "a mutually advantageous conjunction" between personal and work life is exactly what I mean to describe here. Perhaps some phrases—synergy, personalized education, purpose—just need some chicken soup and falsifiability.

[7] For any snooty grammarians that don't like my use of irony here, please offer me a better word and then cram it up your ass.

The following are a set of contexts in which balancers and synergists will experience differing challenges and how to respond.

There isn't a formal onsite office and that absence is particularly problematic for balancers. A vast research literature has documented how brains naturally link spatial context (e.g., work office, bedroom, kitchen, coffee shop, etc.) to memory, affect, emotional intelligence, and more. In the specific context of work, for example, spatial context influences procrastination; if you are trying to work in a room normally meant for play, the work suffers. And the reverse is true: many employees find their homes become work-associated, causing them to neglect personal and family issues. These issues have minimal impact on synergists who already integrate work and personal life, but for balancers, establishing a distinct work environment in the home is crucial. Ideally, everyone would have a dedicated home office, but that is rarely true. So, create a unique workspace that provides a multimodal signal to your brain that defines "working". For example, in a studio apartment, set up a workspace in one corner looking into the room, creating a unique view of your room that you only experience while working. (Don't get lazy and shoot off emails from the futon!)

There is no formalized onsight work schedule anymore, requiring distinct strategies for synergists and balancers. The daily processes of traveling to and from work provided an enormous amount of implicit structure to employees' work lives. For balancers, it gave an unambiguous start and end to the workday. They can largely set aside work responsibilities while away from the office. For synergists, this has always been annoying–all of those emails they send out during the weekend don't get answered. Balancers are less inclined to blur personal and work and inevitably struggle as structures previously embedded in the office environment are stripped away. They need external boundaries in work for successful boundary management at home. They also need the freedom to establish regular availability that mimics a normal workday without social pressures to deviate from it. Without this additional structure they pull back and invest less in their efforts.

Synergists are more adept at micro-transitions between work and personal and don't require the temporal and physical structure offered by a traditional work day. Working remotely can be a good thing for them. Some strategies like conspicuous monitoring may be useful for balancers, but are likely to have negative impacts on synergists. Their problem is that they need to be given more control over their own schedule and flexibility in their engagement or they will burn out. Without change, they are simultaneously answering emails on demand during the traditional workday and engaging extensively outside of work hours. They need the autonomy to set their own flexible schedules, just as balancers need an established structure that defines when they are expected to engage.

Rather than demand all employees behave the same, establish norms that respect different needs.

Working from home comes with unique distractions in the form of leisure activities and familial demands. When employees are given more control over how and when they work it inevitably leads to a greater integration of work and life. For balancers this means more interruptions in both directions. For synergists, this may not be a change and the "interruptions" are part of a natural daily flow. Autonomy for them allows dynamic shifts in priority to deal with tasks, work or life, as they occur.

Generally speaking, more interaction between supervisors and employees decreases family and leisure interruptions of work productivity. Much like conspicuous monitoring, increasing interaction with supervisors drives an implicit demand on employees to remain focused on work. This relationship largely disappears when comparing proactive to more passive personalities. Only more passive employees, those waiting for direction from supervisors, benefit from increased interactions. For balancers, frequent interactions are desirable as long as they are during prescribed periods of time. Without strong meta-learning skills, particularly meta-cognition and emotional intelligence, frequent touch points can be essential for remaining productive. As employees become more proactive in creating work opportunities for themselves, they more naturally integrate work and home together.

All of this has particular implications for working parents. Research shows that working from home increases the number of family interruptions of work as well as the number of work interruptions on family. If you're a balancer, it's the worst of both worlds; if you're a synergist, it's finally the freedom to set your own priorities. And so, just as with distributed work's other challenges, different types of working parents will need different support.

Of course one major difference for many parents is the different working experience of mothers and fathers[8]. Research long before Covid-19 has shown that women do substantially more housework and childcare even when controlling for work hours, but the total lockdown has revealed new frictions for both mothers and fathers. For example, remote work creates greater work-to-family conflict for men; in other words, without clear boundaries, most men increase their contribution to family care less than their wives. Women experience a similar phenomenon, but rather than being driven simply by an increase in remote work hours, it is related to aspects of meta-learning, specifically the inability to disengage from work. Balancer moms (or more generally, all balancers with high conscientiousness) need a work culture that actively supports clear boundaries.

(Covid-19 has created a unique remote work experience for families around the world. Not only are parents working from home, but the children are at home as well, taking classes or simply an unscheduled vacation. This has become the principal source of distraction during the pandemic lockdown. It's not an inherent feature of distributed work that we also need to be full

[8] For same-gender parents the challenges can be different, for example gender wage gaps hit two mom families doubly hard. Single fathers or those from two-father households often report feeling that others judge them more harshly for taking family time. And single moms…well, society has never been kind to you.

time caregivers for our children. Even for synergists, multiple 80-hour-a-week jobs at the same time has been too much.)

Requests for information are a consistent source of distraction for all types of employees. “Do you have the latest numbers…?” “What is the status on…?” “Where can I find the…?” Whether you are a balancer or a synergist, organizations with poor access to information suffer from more interruptions. One of the easiest ways to improve productivity for both groups is to ensure that people don’t need to ask for information trapped in another person’s head. Any organization engaged in distributed work must make all information readily available through collaborative knowledge repositories online. (We’ll discuss asynchronous collaboration in “Rebuilding Culture” and “Distributed Innovation”.)

Complex creative work is suffering in a remote world. Chinese search giant Baidu conducted one of the first analyses of software developer productivity during the pandemic lockdown. They found that productivity increased for simple, modular tasks that could be done by a single developer. Other research has also shown that employees doing complex but non-collaborative tasks were more productive when remote. These findings support the intuition expressed by many that remote work means fewer distractions and greater productivity. But these intuitions are wrong where it counts the most. Baidu’s report shows gains only for routine or solo work–as task scale, complexity, and innovation increased, productivity decreased dramatically.

This early research has strong implications for different types of workers. Balancers will shift towards oversimplified work, steering away from the messiness and blurry boundaries of complex projects. Contrary to intuition, this shift will be greater in more conscientious employees as traditional productivity measures incentivize employees to shift away from complex, collaborative tasks that suffer in remote. For managers, these shifts in balancer productivity will be exacerbated by issues of time, space, and distraction identified above.

Synergists are less likely to shift their work from more complex tasks, as they tend to be more goal-oriented rather than process-oriented. Instead, they will work faster in order to compensate for the inefficiencies of remote work. This might sound great in the short term, but if they feel under constant pressure to deliver and accelerate their work in response, research has shown that creativity decreases, while error rate and chance of burnout increase. While managers should give synergists increased autonomy, this doesn’t mean just throwing problems at them as though they are a boundless source of productivity. More generally, distributed work needs new tools to facilitate complex creative work, particularly asynchronous and semi-asynchronous technologies. We will return to this last point from a company-wide perspective in “Rebuilding Culture”.

For the individual employee, we already know what works in the long-term: meta-learning. Importantly, individual meta-learning “skills” are changeable, and organizations can invest in programs that foster meta-learning throughout their workforce. Over the long run this will reduce mismatch in distributed work, just as we might address hardware and software mismatch in

distributed computing by investing in infrastructure. Actively lifting meta-learning skills and increasing the number of synergists has the wonderful side benefit of increasing job, life, and family satisfaction.

Unfortunately, shifting a workforce towards synergists is an effortful, multi-year process, and it won't include everyone even in the best circumstances. In the short term, companies must deploy differentiated support policies for different types of workers: autonomy for some, conspicuous monitoring for others; explicit boundaries for some, self-regulation for others; and, project complexities that optimally match individual and team capacity. People are different; successful distributed work recognizes this simple truth.

Most of the existing research on remote work has assumed that all employees respond the same to remote work and that one set of policies can support everyone. These assumptions emerged because pre-Covid-19 researchers studied largely homogenous populations, either those that earned the right to work remotely or low-autonomy gig workers. The differences we describe above don't become apparent until everyone is forced out of the office.

So, be honest with yourself and others about what you need. Managers should learn to recognize the clues that something is going wrong, but with self-reflection, you can identify them in yourself as well. Different types of fragility require different types of intervention. For example, the following set of questions from "6 Tips For Managing Remote Employees" can be a helpful framework if you reject the idea that there is one answer for every employee.

"What are the normal working hours for the team?"
Individuals in different time zones can quickly feel isolated and frozen out of distributed work teams if work hours are defined by the convenience of headquarters. Further, defined work hours are really only beneficial to balancers. When answering this question, be flexible.

"How long will it take to get back to each other?"
Predefined expectations can be incredibly helpful for balancers because it provides unambiguous structure to their workday. For synergists, having a set expectation is an attack on their autonomy that undermines their ability to set their own priorities. Everyone needs to be explicit with their answer to this question and also comfortable accepting everyone else's answer.

"How will we notify each other when we will be unavailable and unable to meet these expectations (e.g., out at a doctor's appointment)?"
This is just another distracting request for information and it has the same solution: asynchronous tools.

"Establish a Video-First Culture"
As the author notes, video is powerful because of its ability to convey non-verbal cues. Unfortunately, existing technology is a rather poor medium for non-verbal communication. At

Socos Labs, we have been exploring the possibility of artificial intelligence to enhance non-verbal cues for both individuals and teams, such as conveying the laughter or the boredom of an audience when you can't see every face. Video can only be one part of a solution, however, as it is a real-time (synchronous) tool and research from distributed cognition and distance learning shows that asynchronous tools are essential.

Controlling the Beyond-Control

Within a single computer, computation is principally limited by available memory and processor speed. In distributed computing, the time to complete an operation is dominated by laggy networks. It is crucial to design networks with both minimum latency (time to pass a message) and maximum efficiency (the most information in the fewest, shortest messages). Remote work has the same inherent lagginess. Every synchronous interaction–Zoom meetings, Slack chats, phone calls, or anything else that requires participation in the moment–takes more time than a quick question in the office[9]. Just as in distributed computing, effective distributed work must minimize latencies (wait times for responses) and maximize efficiencies (reducing the number of synchronous interactions).

Distributed cognition makes several observations about the impact of latencies and inefficiencies in collaborative work. In "Rebuilding Culture", we will discuss the large body of research on overcoming latencies in remote education (hint: it's all about asynchronous collaboration and autonomy). In this section, we'll focus on a simple truth: our global experiment in remote work was ill-prepared and involuntary. Many sources of latency are largely outside the control of individual employees. As we've said repeatedly, people are different, and their circumstances are different as well. Unfortunately, research also shows that arbitrary factors, like these context differences, strongly influence our perception of others' quality of work. Given the slightest opportunity, our brains love to make attributions where none exist[10].

There have been many examples of companies successfully adopting limited remote work practice. Xerox, for example, had 11% of its employees in a virtual workforce program. Countries and companies that transition successfully to remote work tend to be large, resource-rich, and have robust digital platforms and payment infrastructure already in place. Additionally, their remote work programs have focused on high-skill professionals that were working in small teams.

---

[9] Remember that even a quick Slack DM includes the time for the recipient to notice the message and the sender to notice the reply. But there is also the attentional cost of leaving one screen and moving to another and then shifting attention back again. These attentional shifts are costly and lack in-person social cues to mediate expectations.

[10] There is a huge research literature on *fundamental attribution errors*. These generally involve attributing some foundational quality to a person based on limited experience, even if we know that the circumstances of the experience were arbitrary or even biased. It's worth noting that this bias shows cultural and individual variability.

What should be obvious is that the experience of 10% of the most highly-paid, highly-educated employees isn't going to generalize well to everyone else. As the rest of the workforce joins them, laggy internet connections, closed schools, under-training, and the simple reality that not all work can be done from home will throw existing best practices into disarray.

Perhaps the most fundamental situational factor is access to high-quality and reliable IT, particularly broadband. Not all companies and countries have invested in adequate IT infrastructure. Poor neighborhoods, rural locations, and historically disadvantaged communities have been left behind in broadband connectivity. Employees working from these locations have poorer network connection when logging in for those crucial video conferences. Although the quality of that connection is both outside their control and not indicative of their quality of work, research suggests that coworkers will slowly but inevitably draw a connection between quality of work and quality of video.

Other IT inequalities exist in remote work that were never issues in the office. Employees might have wildly differing personal hardware, from computers to cameras, and sub-optimal wifi. More subtle but even more crucial, not everyone has a professional workspace available to them. Only the most fortunate will have a dedicated home office, but even the ability to work from the kitchen counter is impacted by crowding in a small household, family and flatmates competing for internet bandwidth, and neighborhood construction[11].

Obviously none of these external factors have anything to do with the qualifications of an employee. In the office, everyone has access to the same infrastructure. At home, someone's postal code can have an outsized influence on their capacity for work. The lagginess of that video connection or the inability of someone to run new software has little to do with an individual employee but negatively impacts our perception of their performance.

For managers, you can combat these biases by providing multiple channels of communication to allow employees to contribute in whatever way is most effective. For example, a meeting centered around Microsoft Teams or Zoom should also leverage simultaneous chat and asynchronous collaborative documents.

For large employers, step up and provide your employees with the infrastructure they need. Supply your workforce with high-quality equipment. Here are some good recommendations: "How to make video calls almost as good as face-to-face". Treat an investment in their home office as an investment in your infrastructure.

In fact, for distributed work to be successful, you must consider community broadband a common good asset to your company, just like a reliable transportation network was for the office. This is not social justice[12]; it's just good business.

---

[11] Particularly in my writing assistant's neighborhood. Apparently she lives in the Tower of Babel.
[12] But why should it be so scary for companies to do the right thing simply because it's right.

In the context of Covid-19, many companies have responded with reduced hours, furloughs, and even layoffs. This only makes sense if you see employees as a fixed asset incapable of change[13], but in fact, it's not only possible for employees to change, that change is crucial for adapting to distributed work. Just as we've discussed upgrading IT infrastructure, how about upgrading your company's human capacity infrastructure. Investing in human capacity is a common good asset, just like transportation and IT; job training pays off even when there is no guarantee that your firm will reap the increased performance benefits of the specific worker you trained. The biggest benefits will not come from "upskilling" for specific hard skills but from investing in developing the kinds of meta-learning–autonomy, communication, self-management–that are most related to success in distributed work.

Investing in training for individual employees is just a start. A number of existing business systems have proven to be incredibly fragile to economic shocks like Covid-19. A more ambitious agenda would conceptually align supply chains with distributed work. Where most of the thinking around remote work has concerned laptop jockeys attending meetings in their pajamas, distributed work like distributed computing is a more generalized idea of shifting away from laggy, fragile, synchronous systems. While we've been talking about employees, this could include entire factories that represent single points of failure in a production line. Even though manufacturing and warehousing require physical locations, these can be broken up into smaller, more dynamic facilities that are less vulnerable to cascading failures or outbreaks. They also have the advantage of being manageable by small teams with flat hierarchies, which as we will see, are drivers of success (and innovation) in distributed work. Distributed work practices themselves might support a shift to distributed manufacturing.

Broadband won't help everyone because not everyone can work remotely. Only "37% of jobs in the United States can be performed entirely at home" and, as we've noted, these jobs mostly focus on professionals in knowledge and creative fields. Manufacturing, logistics and distribution, and a wide variety of other job verticals have traditionally required a physical workforce that can't operate remotely in the same capacity. These industry differences impact certain cities and nations more heavily, as "lower-income economies have a lower share of jobs that can be done at home." The economic disparities between employees that can easily shift to remote work and those that cannot are hard to ignore. Any strategic plan to create a distributed workforce needs to include everyone.

Most of the world has built a work culture around offices, factories, and other physical locations. We're in a moment of profound transition that will expose a great many inequities in employees' ability to work from home. They are still good employees even if we have failed to build an infrastructure designed for remote work. Saez & Zucman argued that we shouldn't let otherwise viable small businesses fail during the lockdown and erode our long-term business capacity. The same applies for human capacity.

---

[13] E.g., this employee is a tool capable only of changing bed sheets in hotel rooms.

Rebuilding Culture

In previous sections we have thought about remote work largely in the context of individual employees and how managers and companies can support them. But a true distributed work strategy needs to think in terms of communities. The goal shouldn't simply be policies that make selected individuals successful, but a culture that supports collaboration and innovation. Supporting a community culture is particularly challenging though, because just as in distributed computing, distributed work suffers from extremely narrow bandwidth that limits communication. Many organizations have tried to overcome this limitation with better technology. While improving our tools can help, the real solution is a culture that is native to distributed collaboration.

In many ways, the bandwidth of communication between two or more people has always been narrow[14]. However, in-person teams are always in implicit contact, whereas remote members are only ever in explicit contact (video meetings or social channels). If we force old habits through narrow pipelines, human communication can only be slowed by working remotely. Some might argue that cutting out unnecessary meetings and distracting socialization is a good thing; as we noted in "Business-as-Usual", however, there is surprisingly little and conflicting evidence that remote work decreases distractions.

Overcoming narrow bandwidth requires increasing the richness of communication as much as possible. There are three prominent themes in the research on remote work culture. The first is the importance of intentionality in establishing that culture. The second is the crucial role of asynchronous communication and collaboration. And the last is the degradation of innovation and inclusion (and some hints on how to overcome it).

From a distributed cognition standpoint, culture is a set of shared tools that improve communication and problem solving. From this perspective, the number one rule of distributed work is don't let chance, laziness, or bad habits define your work culture. For example, scaling effects of online networks cause many companies to form large, undifferentiated communities through their social channels, where smaller, more nimble groups would have existed in offices. I experienced this firsthand while briefing one of the world's largest tech companies. They have many offices around the globe, but when work shifted wholly online, employees outside the headquarters' timezone felt isolated, unappreciated, and excluded. Research has shown that large, centralized teams tend to develop ingroup-outgroup mentalities that promote conflict and degrade coordination. The scale effects of internal social networks can work against you by connecting everyone to everyone else. Organizations must actively establish smaller, more distributed teams.

[14] Mouths are a pretty mediocre technology.

Flat Hierarchies & Shared Leadership

Two crucial factors for the success of small teams are flat hierarchies and shared leadership. Narrow pipelines create communication bottlenecks that become a drag on the productivity of hierarchical teams. When top-down management cultures are constrained by limited communication channels, individual employees often wait for direction in response to unknown or uncertain situations. One reason relatively small, autonomous teams have been most successful in distributed work is that they are not limited by communication bottlenecks.

Unfortunately, in the absence of a strong, established culture, granting autonomy to every employee simply produces chaos. In fact, remote work can even exacerbate these problems because distributed teams over attribute communication failures to incompetence compared with colocated teams. Synergists flourish with autonomy because they possess a diverse set of meta-learning skills–self-assessment, emotional stability, conscientiousness, and resilience. These socioemotional factors provide an internal structure that keeps them aligned with their teams. For balancers, culture is an external structure that plays the same role. It provides them with the explicit boundaries that they need in order to work autonomously. For culture to effectively scaffold all team members, goals, roles, and communication norms must be formalized from the very beginning.

Additional research on remote work reveals that people's relationships with their colleagues suffer compared to in-office work (though anecdotal evidence suggests a decreased number of downward comparisons, i.e. people worry less about others slacking off). A 2013 Gallup poll found that remote workers log an extra four hours per week on average; these increased hours result in additional work stress as noted in "People Are Different" and "Controlling the Beyond-Control". Much like the misattribution of communication failures, this additional stress can also be misattributed to remote colleagues and lead to the degradation of coworker relationships. This will be particularly prominent for balancers in teams that lack explicit boundary norms.

Research on successful teams, both remote and in-office, support the importance of all the findings above. For example, Google's well known internal research indicates that psychological safety is a principal predictor of team success. This meant that team members had a set of shared norms and trusted one another without need for constant communication. In fact, members of Google's most successful teams had fewer synchronous communications internally and spent more time working autonomously. A study by the company Dropbox on collaborations between successful academic teams revealed that the most successful teams were small, autonomous, and maintained flat hierarchies. These qualities agree with the research findings noted previously, but their analysis also suggests the importance of role-modeling by senior members. Other research indicates that sharing stories company-wide can play the role of role-modeling in distributed work. Stories of successful problem solving by teams are the principal mechanism of establishing cultural norms. All of this research suggests the importance

of transparent role-modeling in establishing culture within organizations, particularly in distributed work where potential role-models would otherwise be hidden.

From my own research on over 60,000 companies, it's clear that culture is not the slogans a company writes on its walls, but rather what its employees actually do. Inspired by an analysis of personality characteristics conducted by Facebook, I collected data on tens of millions of employees at these companies and analyzed their own self-descriptions as a way to understand the "personality" of a company. This machine-learning analysis revealed 186 dimensions of cultural variability across companies. Some of the most prominent were, for example, management- vs. employee-driven[15] cultures and research- vs. process-driven cultures. Engineers at management-driven organizations bragged about how they would return beautiful products when given explicit instructions; engineers at employee-driven companies bragged about themselves and how they would solve problems. Unsurprisingly, innovation correlated strongly with employee-driven cultures, while management-driven was related to better risk management[16]. As for process-driven cultures, let's just say they were very common in the DC Beltway.

One of the most important takeaways from my analysis of company cultures is that they are the product of the employees, not the mandate of management. Because culture in my model was derived from employee behavior, I was able to experiment with predicting cultural fit simply by looking at how well a potential employee fit the model of a given company. One of the most important takeaways from this insight is that establishing (or re-establishing) a culture doesn't come without effort. It requires shifting the norms of an entire population of employees. Driving such a shift requires at least two components: an intentional framework of cultural norms and role-modeling of those norms by leadership. The single biggest impact leaders have on their teams is not their memos or strategic plans, but their actual actions, the sacrifices they make to accomplish organizational goals. Those actions are often hidden in the best of times, and in remote work, entirely buried. It's important these stories be shared transparently.

Another finding from my work is that the best teams are intentionally based on *complementary diversity*. This means that teams are designed so that members' relative strengths and weaknesses are complementary to one another. These could be differences in expertise, experience, personality, identity, or many other dimensions. Composing teams based on complementary diversity allows peer role-modeling to play a powerful role in developing team capacity, but new teams by their very nature lack psychological safety. People that are different tend to speak a different language, at least metaphorically, and sometimes literally. Strong

---

[15] I later discovered that this is a well-known cultural dimension in sociology, traditionally referred to as *tightness vs. looseness*. Tight cultures tend to have less crime, teen pregnancy, and greater stability; loose cultures invent more stuff and apparently have a lot more fun.

[16] Hierarchical cultures don't develop out of nowhere. While they are maladapted to distributed work, hierarchies inevitably develop in any organization threatened by systemic risk (i.e., banks or militaries). Some poor bastard needs to take the blame.

cultural norms are crucial for combating the stress and misattributions that tend to dominate new, diverse teams.

While intentionality and role-modeling are essential for establishing cultural norms, there are a number of well-grounded tactics that can help overcome narrow pipelines and improve communication efficiencies immediately. One of the most classic is to set a specific agenda with desired outcomes *pre-identified* for all meetings. This is as true in distributed work as it is in the office. For a more flexible approach to meetings, some have suggested "borrowing the idea of 'office hours' from academia". Rather than having a fixed meeting, managers and employees more generally can post links to weekly video sessions where they can respond synchronously to any small questions that have popped up throughout the week. The idea is that this would cut down on the number of distracting interruptions by having a fixed and reliable time in which anyone can ask a question. (Having hosted a great many office hours myself, I can say that it can be an enjoyable distraction to simply have a free hour where you've dedicated yourself to a mix of answers and conversations with a random assortment of characters.)

Asynchronous

Another common recommendation is the use of message boards and chats as a solution for communicating while teleworking. Unfortunately, research hasn't been as kind to this idea. As noted above, undifferentiated message boards and internal chat technologies can tend to become rather bloated and create in-group/out-group relationships in a company. They are also subject to *network capture*–undifferentiated social interactions that tend to have a winner-take-all quality–in which a small number of individuals dominate the conversation. While a thousand separate conversations within small groups of people may be a poor way to disseminate information, one massive conversation between five thousand people tends to erode innovation and (ironically) inclusion.

Meetings and chats are explicitly synchronous communication technologies. Email lives somewhere in between. Many people use it almost as a kind of chat, shooting off quick questions and expecting quick responses. For others, email is like letter writing in the 19th century, a labored, near-literary process[17]. A number of studies of email use have shown that more time spent on email increases stress and decreases perceived productivity. It turns out that most of this subjective productivity loss comes from people who allow emails to interrupt their workflow. Psychological literature suggests that this might be a form of learned helplessness, in which intrusive events outside an individual's control decrease their long-term performance at a neural level on an unrelated task. Instead, I recommend treating email as an asynchronous technology. Rather than immediately responding to incoming emails, individuals

[17] You might be able to guess where I fall on this spectrum from my rather agonized use of semicolons in tweets. Yours truly, Vivienne L'Ecuyer Ming, PhD.

need both the team norms and the self-regulation to address emails in batches when it fits with their workflow. To paraphrase the CEO of Automattic, I'll get to it when it's the right thing to do[18].

Used in this way, email shifts towards asynchronous communication. In distributed computing, "**asynchronous**" is a challenge to be overcome. Most distributed systems operate on queues, where pieces of information and tasks are held in the order received. Individual computers have established rules of how to handle and order the queue, what to do when it overflows, and how to interact with other computers on their own schedules.

In distributed cognition, "asynchronous" is a superpower. It allows teams to make collective progress on projects even when they're not in direct synchronous communication. Intelligent systems (like us (well...most of the time[19])) have an advantage over the simple queuing algorithms of distributed computing in that we can build mental models of one another. We can infer the conditions and intentions of all of our teammates and adapt our actions dynamically, whereas hard-coded systems must cover all possible outcomes or break. As such, distributed work is much more robust than distributed computing when the three key ingredients are present: autonomy, psychological safety, and norms.

Many of those key communication norms are very grounded and tactical. For example, teams must enforce norms about using modes of communication that match the urgency of the message. Don't send an email if you absolutely need an immediate response; don't send a text message about a complex and emotional question. To make the most effective use of the asynchronous strength of email, we can actually follow the norms we identified above for meetings. Every email should have sufficient background, required outcomes, and an explicit due date. This gives the recipient the information they need to make their own decision about the contents of the email.

While email isn't inherently asynchronous[20], more and more digital tools are designed to be natively collaborative. In fact, it's gotten to the point where I am disappointed in cloud software if two people can't work on the same document at the same time**.** Unlike discussion forums and chats, tools such as Google Docs, git repos, Notion, and wikis aren't just a log of every debate or offhand comment that led to a finished project. They are more than just a record of a process; their true value emerges from the evolving synthesis of all the collaborative learning that got you there.

Perhaps the most famous asynchronous tool of them all is Wikipedia, the world's largest encyclopedia[21]. Each hyperlinked article might represent years of learning and negotiation between the editors collaborating on that subject. While the wiki tools log all of the individual

---

[18] Actually, the original quotation was "I'll get to it when it suits me," reflecting the exact kind of sterling personality that has given tech bros the reputation that they have in the world today.
[19] Well...some of the time.
[20] Don't blame me–I was actually a big fan of Google Wave.
[21] And if you didn't already know that, I cannot help you.

edits, contributions, and comments, the live article represents a synthesis of learning, creation, and social communication. The wiki creation process doesn't force Wikipedia editors to be in the same place at the same time, but it does transform their raw communication into shared knowledge. (Could you imagine if Wikipedia was a series of Slack channels detailing all of the arguments and debates over what constitutes the "truth" on a given subject?)

With the collaboration and transformation inherent to asynchronous tools, every evolving project document represents the culture underlying the work, or as Thomas Malone would term it, a "supermind". The dynamics of asynchronous tools capture the intangible capital of a team and begin to break down the distinction between planning and working. They force an explicit embodiment of the collaboration–not just a list of action items but a shared understanding.

Asynchronous collaboration is where distributed work will shine, but it comes with unique challenges. I've led hundreds of projects across academia, industry, and the arts and learned, rather painfully, the value of asynchronous work. These days I use a combination of Google Docs and Notion (though I'm always on the hunt for new tools) to make the act of collaboration tangible. One of the most important rules I've learned is, "It must be in the doc." Side notes and forks[22] are absolutely not allowed; it is the collaborative resolution process that creates the value add of asynchronous tools, and so all work on a project must be subject to this group computation. For the same reason, it's imperative that contributors aren't passive, waiting to be told their role in some resulting list of action items; ideal asychnolous tools directly touch the finished product. Yet, these tools are also a form of communication. A useful framing is to consider whether a completely new team member could take up your role based only on the tools. They should be able to walk into the project and rapidly join that supermind.

This doesn't mean there's no role for "meetings" in distributed work–it can and should be a hybrid of asynchronous and synchronous. Research demonstrates that remote meetings are more effective when preceded by asynchronous discussions. The findings become even more persuasive when we look at the decades of research on remote learning in education.

Remote learning has been an active field of education research since the 80s and 90s. In recent years with the rise of platforms such as Khan Academy, Coursera, and edX, we've begun to understand how to create learning experiences that are native to remote technology, rather than simply being filmed lectures. Asynchronous learning is one of the core features of *computer-supported collaborative learning*.

Asynchronous collaborative learning involves a mix of tools: pre-recorded lectures for self-paced viewing, collaborative wikis, discussion forums, and online learning environments where students can collaborate on assigned problems such as virtual whiteboards and lab spaces.

---

[22] By this I mean, "I made a separate copy of the document with my changes," without the effort to resolve those changes back into the document. Forking in coding is common to prevent breaking the system, but it is always resolved back into the head. Code also attempts to be modularized in ways that other forms of distributed work are not.

These technologies free students from being yoked to one another's pace, but more importantly transform learning into a student-centered, active-learning experience. Instead of students passively listening to lectures together and completing synchronous in-class work, instructors use tools to monitor student engagement and progress on disparate problems, giving support for students' individual experiences as needed. The main challenge is driving engagement with the students and preventing dropout. This is particularly important for students who fall to the periphery of social learning networks.

One study found that while traditional lectures yielded better student performance in fully synchronous classroom settings, when classes go online, asynchronous, problem-based learning produced the best results. The change in the medium from in-person to online qualitatively changed the process of learning. Another study showed, online students that participate in asynchronous discussions perform better than those that only follow class lectures and other traditional teaching practices, and still another found, "...asynchronous peer-to-peer discussion is more effective than traditional classroom lecture-discussion for undergraduate students."

The research on successful asynchronous learning reveals a shift from a passive teacher-centric to an active student-centric model. Beyond the specific details of asynchronous tools, this dichotomy mirrors the management-driven vs. employee-driven (tightness/looseness) cultural dimension that I have previously identified. Remote learning flipped the hierarchy of the classroom, with teachers supporting individual students rather than leading the class. Remote work has turned traditional work hierarchies on their head as well, with leaders at the bottom acting as support for their employees rather than delegating from the top. It should stay that way.

All of this tells us what distributed work should look like. Distributed managers and instructors are the glue that keeps teams working towards a goal by establishing and role modeling norms through their own creative contributions and preventing dropout, burnout, and disengagement. It's a shift from one person acting as a bottleneck of information and direction. Managing asynchronous teams also means knowing how to manage asynchronous technologies–it is an integral part of their role in the team. They are the ones that will set the norm for how an entire team uses tools. If they are lazy in their use of the technology, everyone else will follow suit.

Part of the beauty of asynchronous collaboration is its ability to bring balancers and synergists together on their own terms. The tools allow for flexibility in the norms of interaction (though these must still be explicitly agreed upon ahead of time). "I can only be contacted in these hours." That's okay if during that time you are engaged with your team's asynchronous docs and wikis, so that I can contribute on my schedule. "You can contact me whenever, but I will get back to you when I have time, so don't bug me," is also transformed by adopting asynchronous tools. With few demands for synchronous interaction, individuals can engage in the way that best integrates with the rest of their lives.

Distributed Innovation

“The meeting of two personalities is like the contact of two chemical substances: if there is any reaction, both are transformed.” Jung’s simile is perhaps my all time favorite metaphor. On its surface, Brownian people drift through life, occasionally colliding and, by chance, transform each other. Jung’s interpersonal chemistry captures two of the greatest worries about remote work: without those chance collisions what drives innovation, and without interpersonal transformation how do we increase inclusion? Over the last year, I’ve been asked these two questions about remote work more than any other[23].

If innovation is just an accumulation of randomly colliding ideas and personalities, as so many assume, all we can do is bring more and more people together, hoping to scale up the collisions. Density brings serendipity. Innovation requires bigger cities, more universities, in-person conferences, and a seemingly absurd amount of time spent lingering around the watercooler waiting for just the right off-hand comment to spark an “ah-ha” moment.

It’s not an unreasonable metaphor. We already know that research, innovation, and industry are concentrated disproportionately in the largest cities. Just ten US cities account for 48% of its patents and 33% of its GDP, even though they make up less than a quarter of the population. Across science and industry, as activities increase in complexity, they also increase in spatial concentration. Some have even argued that a single parameter, population, accounts for all productivity and innovation variation across cities[24]. If spatial concentration and pure serendipity are crucial to innovation, then “remote innovation” might quickly displace “army intelligence” atop the oxymoron rankings.

Insights from distributed computing don’t offer much hope. In these distributed systems, all communication is passed through narrow pipes, requiring potentially complex (de)serialization algorithms and long latencies. When all work passes through these same pipes, the sort of scale-dependent serendipity assumed to drive innovation is nearly impossible. Virtual watercoolers lose the chance interactions of the real thing, as everyone is there for the explicit purpose of socializing. Company-wide open blogs and social channels are popular remedies for isolation, but as we’ve noted earlier, winner-take-all effects in large-scale networks can easily reduce diversity of thought rather than promote innovation.

---

[23] Yes, I am the kind of asshole that waits until the closing sections of an article to answer the most important questions. But we had some good times along the way, right?

[24] Except that the author needs a separate parameter for each region of the world…which sounds like many parameters to me. More importantly for our purposes, it sounds like something about the architecture of the social graph rather than simply its density is important to innovation.

If we're just particles waiting for a chance reaction[25] through rate-limiting pipes, then it's hardly surprising the big, complex projects at Baidu ground to a halt when the world went remote. Worse still, all of this is within the broader decline in VC-backed early-stage startups and corporate innovation common during recessions. Capital flows away from the truly novel to sure up more mature bets.

Fortunately, Jung's metaphor has a deeper mapping between the complexities of chemical reactions and interpersonal relationships. Imagine the sophistication of advanced chemical engineering, or even coronavirus vaccine development, reduced to the brutish simplicity of two enormous vats of chemicals dumped together in a tub. No catalysts or buffering agents. No stoichiometry. No knowledge of the underlying chemistry at all. Just swirling molecules at massive scale.

While the interaction of any two specific molecules may be a fundamentally probabilistic process, modern chemistry grants detailed control of its product by biasing the statistics towards the outcomes we desire. Perhaps the pipes carrying distributed work can play the role of pipettes carrying chemicals and transform innovation itself.

Innovating on Innovation

Despite the proliferation of "me too" startup incubators and dubious corporate innovation handbooks, no "innovation engineering" fuels creativity the way chemical engineering fuels Pfizer, Dow, and Chevron. That doesn't mean we know nothing. For decades, researchers have been quantifying the neuroscience, psychology, and social dynamics of innovation. For example, we've identified the complex and effortful interplay of brain networks that drive our capacity to generate original ideas, and have even identified the brain regions behind high-quality jazz improvisation. We've identified the subtle interplay of attention, memory, and reward that leads to those almost mythical ah-ha!-moments, although it seems eureka is much less bolt-of-lightning than working memory and fluid reasoning[26]. And the brains of Nobel prize winners in chemistry and literature look very similar in moments of creativity. Perhaps it's not so surprising, then, that Nobel prize winning scientists are 3 times as likely as the general public to pursue artistic interests outside the lab, but importantly, their art and their science complement each other, reflecting the same underlying passions.

Most crucial to remote work, distributed cognition research has revealed much more about collaborative innovation than just smashing people together. Consistent with the basic idea that density drives collaborative innovation, two labs that are close to each other on university campuses are 3.5 times as likely to collaborate, but the story immediately becomes more

---

[25] Possibly a love reaction–I, for one, am certainly sick of sitting 'round here trying to write this book. Is all innovation just dancing in the dark?

[26] There is even research looking at the physiological correlates of innovation in which "participants who presented new ideas demonstrated higher levels of galvanic skin response, indicative of engagement, emotional arousal or cognitive load."

complex. For example, once a collaboration has begun, separation doesn't disrupt the relationship. In fact, separation improves the quality of research outcomes by promoting more independent "research trajectories". The original relationship established by colocation is the catalyst for collaborative innovation, but without a rate limiting reagent to prevent cognitive homogenization, the value of that relationship reaches premature equilibrium with little more to be gained[27].

The research on how colocation affects innovation reveals a much more complex story than *density=serendipity*. While colocated labs are much more likely to collaborate, the novelty and quality of the projects they produce drop off. This tradeoff clearly indicates that while we want people to interact, simply pouring everyone into the same vat is as brutish and inefficient as one might imagine. Rather than an undifferentiated equilibrium, our goal for innovation must be dynamic allostasis. The pipes connecting us together can be an advantage, titrating our interactions to optimize the tradeoff.

For example, the flow of knowledge carried by business travelers from one city causally changes the growth in economic activity in a destination city. This finding shows that if you change those travel networks–snipping a connection here, adding another there–you change the city itself. Innovation and economic activity within big cities are about much more than size or density. Like personalities, the mixing of cities leaves both transformed.

Optimizing innovation is not just a matter of how we structure people but how we structure knowledge and information. Searching for books, articles, and websites would be nearly impossible without some kind of hierarchical information structure. Imagine walking into a library [28] where the books are simply strewn around the floor with no structure–no French literature section or "Children's Corner". You have to kick *Madame Bovary* out of the way to find *The Lorax*. But if you're not trying to find a specific book, if what you want is a novel idea, then too much structure also holds you back. When information is presented in a flat, unorganized structure it promotes cognitive flexibility, improving creative problem solving. Flat information structures even helped individuals come up with more creative career ideas.

There is a parallel between research on knowledge structures and spatial structures in promoting innovation. On their surface, both fundamentally rely on dumb luck powered by masses of unstructured interaction. But just as titrating interactions between collaborators might maximize innovation, dynamic information structures might act as catalysts for new insights. For example, participants were more likely to develop useful insights in creative problem solving experiments when parallels to seemingly unrelated problems were made salient.

Another set of experiments in complex problem solving showed that storing subjects' best past solutions for speedy retrieval increased mean performance but decreased exploration.

---

[27] I refuse to apologize for beating this metaphor to death, and if you're sick of it already, oh my goodness...

[28] A "library" is a prehistoric structure in which ancient humans stored all information not worth paying for.

Fascinatingly, this research showed that ubiquitous access to past ideas had the same effect as unrestricted social influence. When people are allowed to interact without restriction while exploring complex problems, the average problem-solving performance of the group increases; however, hidden by this average performance increase is a *drop in the quality of the best ideas*. *Satisficing innovations*[29] both crowd out better, later innovations and preempt continued exploration. The researchers found that the best solutions emerged when a period of separation preceded social interaction. In both knowledge structures and social networks, something as simple as making access intermittent can act as a rate limiting reagent, improving average performance while maintaining the best solutions.

The structure of interaction within a team might even be more important than interactions across a network. A study of collaboration within academic labs reveals the crucial importance of the internal social architecture in tightly connected teams. The researchers were able to leverage data from the asynchronous tool Dropbox to track interactions between collaborators in fine detail. They found that elite teams tended to be small, possibly counteracting the homogenizing effect of large groups. They had flat hierarchies, with group members contributing to projects more equitably. While their contributions were equitable, the best teams also evidenced complimentary diversity: individuals' contributions reflected their domains of expertise and differences in personal experience. But even elite teams weren't perfect. The research found that elite labs were more likely than others to collaborate based on preexisting relationships. In following the path of least resistance, elite labs fail to capture the innovation capacity of novel collaborations.

In 2012, Google launched Project Aristotle to understand why certain teams were more successful than others. On the surface, highly successful teams could be wildly different, but some teams leveraged those differences in powerful ways that others could not. Strong norms within a team seemed to be a crucial factor for success if those norms supported two specific qualities. First, just as with the Dropbox study, elite teams had flatter hierarchies with more equitable and complementary contributions from individual members. In an echo of a large body of research on creative problem solving, when a lone member dominated, collective intelligence dropped. Second, elite teams demonstrated strong average perspective taking. Strong norms promoting prosocial perspective taking in turn promotes psychological safety, a form of interpersonal trust that rewards constructive risk-taking. In the absence of psychological safety, individual members withhold their contributions and the quality of both the average and best solutions drop. Psychological safety buffers teams from large variations in personalities and identity, turning an acid comment into a creative catalyst.

We've only touched on a tiny portion of the vast research literature on innovation and creative problem solving, and yet already patterns emerge. There is a microstructure within individual innovation cells (e.g., flat hierarchies within teams). These individual cells interact through a

[29] Satisficing: halting at the first satisfactory solution that one comes across; the tendency of innovators to accept "good-enough" solutions rather than continuing to look for the optimal solution.

macrostructure of inter-team collaborations and information flows. For innovation engineering to emerge from the existing science, we must grow our understanding of the macrostructure beyond the passive world of density=serendipity. Those narrow pipes can give us fine control of the social graph of innovation, but what do we do with them? One candidate model of innovation engineering comes from the Defense Advanced Research Projects Agency (DARPA).

Like all major US funding agencies (e.g., NIH and NSF), DARPA exists to promote research and innovation principally by controlling the purse strings of science. Unlike its peer institutions, DARPA is both small and flat. The individual DARPA program managers (PMs) do more than decide who gets funded. They leverage an unusual amount of flexibility and autonomy to drive innovation within a specific domain by bringing together disparate cells, whether academic teams or industry labs, for workshops and collaborations. In essence, they are innovation matchmakers, peeking in on the activities of individual cells and strategically manipulating connections between them. In a review of organizational innovation, a group of researchers identify the existence of "idea scouts" and "idea connectors"[30] that have a disproportionate impact on innovation; in many ways, the DARPA PM formalizes both of these roles in a single individual that combines the technical insight of scouts and the social engineering of idea connectors.

The autonomy of DARPA PMs creates space between funding oversight and risky projects. They de-risk the work of individual cells, allowing them to explore ideas that wouldn't normally get funded. Of course, that also means that many projects "fail", but the matchmaker PM can share the insight and learning from those failures via their network. In this way, the PMs themselves replace density=serendipity as the macrostructure of innovation, connecting the microstructure of individual cells to capture the full capacity of the system.

The distinction between macro- and microstructure takes Jung's metaphor far from a vat of undifferentiated chemicals and into the realm of biochemistry, where "inside/outside" dominates over density. Membranes allow cells to control their interaction across the intercellular matrix while preserving the integrity of their internal processes. Individual cells are able to interact, but those interactions are modulated by their membranes. While narrow pipes are a drag on so much of distributed work, in innovation engineering they can be transmembrane protein, ion channel, and extracellular matrix all rolled into one.

Remote innovation might seem like an oxymoron, but in fact distributed work might be the crucial catalyst for innovation engineering to finally emerge. We know its microstructure: flat, small, equitable cells leveraging complimentary diversity and psychological safety. We begin to understand its macrostructure: allostatic tensions and finely tuned networks. Distributed work gives us both unique insight into the microstructure and unique control of the macrostructure.

---

[30] Other research suggests that an explicit "connector" role is not needed if there are other means of information diffusion. We'll explore the forms of those other means below.

Innovation Engineering

So, let's invent innovation engineering. Fortunately, we're not starting from scratch. At Socos Labs, I've been exploring these ideas for years, long before Covid-19 forced us all to work from home. Our innovation engineering framework incorporates tools to analyze the microstructure dynamics of ideas and personalities within each cell and mechanisms to control the macrostructure of collaboration and knowledge dissemination between cells. The basic framework is illustrated in Figure 1. It begins with *incubation*, where individuals cycle between isolation and collaboration within their cell. Next is a novel but crucial process of *maturation*, in which innovation cells collaborate together in small clusters to refine an idea into a robust practice. Finally, clusters engage in strategic *dissemination* to establish new practices across the community.

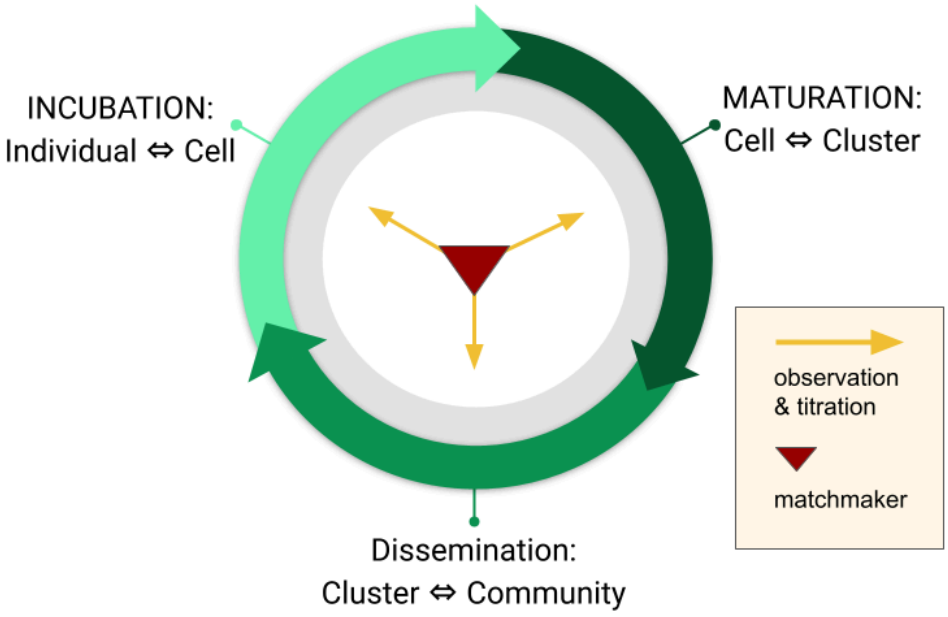

Fig. 1: Innovation engineering framework.

At the center of this framework is the matchmaker, a new role inside organizations modeled on DARPA's PMs and "idea scouts". In the following section, we'll describe how well-trained matchmakers track incubation of novel ideas and dynamically titrate the connections between individuals, cells, and clusters to accelerate innovation. Our illustration of this process uses examples of machine learning tools that augment matchmaker capabilities, but Socos's

framework isn't contingent on technology–matchmakers should be able to carry out these steps independent of any specific tool[31].

Matchmakers peek inside each cell's microstructure to observe the internal dynamics of creation and communication. While this occasionally means joining in on synchronous conversations[32], asynchronous tools give matchmakers access to the state of innovation within and across cells. Socos Labs and others have developed tools to help matchmakers manage the allostatic balance between individual and collaborative ideation. Given the importance of interpersonal dynamics to innovation (e.g., complementary diversity), our framework leverages personality modeling, critical for the mixing and remixing of collaborators. We've also developed models of the conceptual space of new ideas, giving matchmakers greater insight into the landscape of innovation. Finally, our framework models the temporal dynamics of innovation, allowing matchmakers to form new connections between cells and ideas at just the right moment, or even slow the flow of intercell communication when novelty is being drowned out.

Figure 2 captures some of the dynamics of the incubation process within cells (broad purple circles). In the figure, the matchmaker is shown outside of any cell but directly monitoring activity across asynchronous tools and knowledge resources. A cell of six densely connected individuals is shown collaborating around a central asynchronous document, forming a supermind. Individually, they may be engaged in isolated ideation (person 1), in synchronous conversation (4, 5, and 6), or in information gathering outside the cell, either with knowledge resources (2) or other people (4).

[31] Our previous research shows that the greatest productivity gains from AI come from augmenting individuals who combine domain expertise with meta-learning. We call this *creative complementarity*. This finding indicates that the best matchmakers must be highly capable even without AI tools, even as the most effective systems leverage both.

[32] Matchmakers must be fastidious about not intruding on the process of a given cell and becoming an untracked vector of information diffusion. Fortunately, asynchronous tools still allow them to be a part of each supermind.

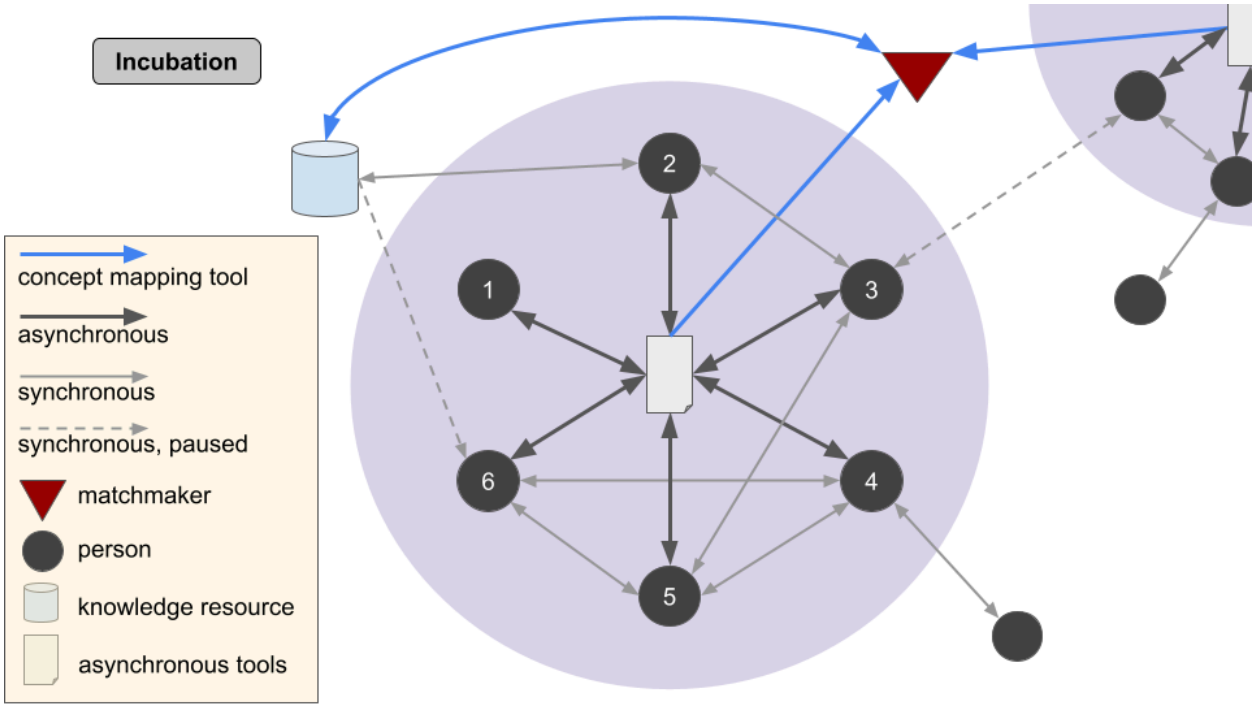

Fig. 2: Incubation and the microstructure of innovation cells.

These six individuals aren't brought together arbitrarily. Some might be long-term collaborators, while others are joining for the first time to bring unique expertise or perspectives. Matchmakers and project managers work together to design complimentary diversity into cells, exploring mixtures of personalities that maximize innovation. Personality modeling has been used to recommend research collaborations, predict team performance, and study peer-to-peer interactions in education. My first experience in this domain was a project to identify beneficial student cohorts in massive online courses. Unsurprisingly given everything we've learned, dumping 30,000 students into an unstructured discussion forum is not a recipe for success. Grouping students into smaller cohorts of 5-10 peers with their own subnetwork improved course outcomes compared to students without a cohort. But we also found some of the first evidence of complimentary diversity, strong indications that a successful cohort needed the right mix of backgrounds, personalities, and experiences. There wasn't a single pattern for success, but active cohort matching improved outcomes[33] by 34% above random cohorts alone.

In a completely different domain, I developed an algorithm to recommend novel connections at professional conferences. The algorithm worked by analyzing the personal and social profiles of pairs of individuals (dyads). Importantly, we discovered that the most successful connections were not defined by similarity alone (homophily). While most people shared touchpoints of similarity, likely facilitating psychological safety, relationship patterns across dyads were predominantly about differences, but those differences weren't arbitrary–everyone had patterns.

[33] This was a competency-based program, without grades; students remained in the course until satisfying all of the competency requirements. Performance was measured as a combination of time to completion and drop-out rates.

For example, one pattern emerged that actually applied to me. We might call it the “geeks hang with nerds” pattern: while I’m a highfalutin scientist that’s rather snooty about my science fiction and doesn’t fully understand why anyone would attend Comic Con, a dominant pattern in my social network are Comic Con fangirl engineers with an appreciation for the campy.

Based on these past projects, our framework models personalities using a combination of explicit and implicit measures (e.g., smart survey methods and language modeling across company communications)[34]. The insights derived from this system are crucial for maintaining the critical allostatic tension between trust and novelty. As we’ve described above, as cells work together over time, the novelty and impact of their innovations decrease. Our framework gives matchmakers the ability to subtly mix and remix collaborations, keeping them at their critical point. This remixing infuses existing cells with both new knowledge and new interpersonal dynamics, but it depends on matchmakers identifying the right pairing. (As we will discuss in the next section, there is a paradoxical tension at the heart of successful innovation collaborations that all organizations must overcome.)

While personalities are crucial, innovations are the actual product. Innovative ideas can be visualized as a dynamic map of concepts, with neighboring ideas having greater influence on each other than those further away. In our framework, matchmakers can use these maps to nudge apart cells to prevent herding around obvious solutions and connect seemingly disparate ideas to uncover novel synergies. Our tool emerged out of an internal project, Cognitive Overflow, which I developed years ago to help map my labyrinthine library of research notes. I would frequently find myself writing on a subject with vague memories of a relevant research article or previous project. After an hour of fruitless keyword searches in my notes or online, I’d have nothing to show for the effort but tears and curses. Cognitive Overflow created a personalized conceptual map of my library. As I wrote, it pulled up relevant research that I could accept or reject without ever running a search. I could even use Cognitive Overflow to remap my library to reflect other people’s conceptual maps and uncover novel connections that might not have occurred to me naturally.

In our framework, this tool can be used within cells to accelerate knowledge discovery beyond either flat or static knowledge structures; however, it plays a distinct role for matchmakers. By applying similar concept maps to the evolving content captured by asynchronous tools[35] (blue lines in Fig. 2), matchmakers can track and even guide the dynamics of idea creation. They can nudge individuals away from less productive concept spaces or from herding around low-risk

[34] Note that there are profound ethical implications and opportunities for abuse in modeling personalities within organizations. Companies have been known to use information like this to prevent unionization or justify layoffs. In our work we are exploring what’s possible, but any real world application would need substantial safeguards to prevent almost inevitable abuse.

[35] If innovation is more than density=serendipity one might wonder why no one was doing this before. Apart from the obvious answer that we are all phenomenally lazy in our own ways and that managing innovative people is the truest experience of herding cats, we lacked the tools to make innovation engineering effective.

ideas, spreading bets and encouraging greater exploration[36]. In Figure 2, the matchmaker is nudging person 2 with a novel knowledge structure (like Cognitive Overflow's map sharing) to make underexplored concepts more salient. Access can also be paused, as with person 6, to prevent exposure to satisficing solutions that might head off future exploration. Even more powerful nudges come from reconfiguring the social connections between cells, such as how the matchmaker has paused the social connection between person 3 and an individual in a neighboring cell.

Crucially, matchmakers' nudges must come at just the right time. For example, an extensive research literature has identified the importance of incubation periods for innovation. Neuroscience has revealed that incubation periods involve unconscious processing of the problem rather than mere distractions or interruptions. Cells that engage in independent work before group collaboration generate more ideas of higher quality and come away with a deeper understanding of the problem. In fact, if matchmakers bring collaborators together too soon, the number and quality of new ideas decreases.

I had my own rather exaggerated version of this experience during my first year of graduate school. I had "invented" a novel algorithm for decomposing sounds and images into a highly efficient "dictionaries" of "code words"[37]. I had convergence proofs and everything. After writing a couple papers and presenting this as my first year project, my research advisor said, "You might want to read this," and handed me Stephen Mallat's paper on *matching pursuit.* It turns out that I had re-invented an algorithm that had already been invented by the guy that had literally written the book on wavelets[38]. I could have felt bad about discovering that I hadn't discovered anything[39], but in fact my understanding of the material and its application to the brain was greatly enhanced by all the work that I had done on my own. When I later learned that the algorithm was known by mathematicians as von Neumann's method, I decided to just be proud that I had the same idea as a man who's in the running for smartest person of all time. And as it turns out, I was able to stand on the shoulders of these giants and convert the algorithm to a truly novel form of machine learning that could learn the optimal dictionary directly from the sounds and images themselves.

I've always wondered whether my advisor intentionally withheld Mallat's research to spur my own creativity or if he just didn't see the connection until we started writing our paper[40]. Whether

---

[36] By the way, I've had a long-held vision of teachers doing the same thing to encourage exploration and innovative thinking in the classroom.

[37] If that means nothing to you, congratulations!–you are a normal human being. If on the other hand, you'd like to learn more, then you're a filthy, filthy geek and you can read up in my paper "Efficient Auditory Coding".

[38] Again, your incomprehension or curiosity is diagnostic of so much about you in this moment.

[39] But seriously, Mike, did you have to wait an entire year before letting on that I was not the first person to ply these waters?

[40] I loved working with Mike, but I was also his first student to complete a PhD, and I think he was still figuring out this whole "advisor" thing. He once asked me if I was attending an elite private conference

intentional or not, he essentially played the role of matchmaker, blocking a macroscale connection which would have prevented incubation of my own ideas. If I'd had the "answer" from the textbook I would never have fully explored the concepts that led to my breakthrough. Research supports this seemingly paradoxical intuition: slowing down the process of knowledge diffusion in dynamic networks actually speeds up innovation. Further, innovation spreads more slowly through densely-connected networks, like a company's social channels, and while highly efficient networks can quickly identify satisficing solutions to complex problems, collective intelligence (long-run innovation) suffers. Both poorly connected and densely connected communities perform worse than moderately connected. This inverse relationship between knowledge diffusion and innovation diffusion may seem confusing if you assume an innovation is just a piece of knowledge. Instead, these findings illuminate a profound distinction: knowledge is information; innovation is practice.

I have a long and colorful history of inventing broccoli. I have detailed analyses on its nutritional value and extensive experiments proving its health effects. Every time, I'm convinced my broccoli will revolutionize the market and people will beg to buy it. On sales calls, customers love the idea of broccoli–they're blown away by all of my scientific research on this weird, green, alien-looking product, but then they say, "So...what the hell am I supposed to do with it?"

My potential customers tell me they're committed to their brownies. The issue isn't just that the brownies taste better; brownies have established distribution channels, tax incentives, and best practices. They "get" brownies.

They believe my research about broccoli, but nobody has any experience with it actually improving their lives. It all seems so abstract. Then they take a bite of raw broccoli, and I can see that the sale is lost. Neither they nor I have any recipes for broccoli. I assumed that the idea of broccoli, with all of my scientific evidence, would be enough to spark my customers' creativity, but nobody is willing to put in the effort to come up with their own recipe just so they can buy my product. Everybody loves the idea of broccoli, but nobody wants to buy it[41].

Over the years, I've invented broccoli across numerous industries and scientific disciplines. Again and again I've looked on as people have said, "This is the most amazing thing I've ever seen, but what the hell am I supposed to do with it?" My new ideas, even fully developed new technologies, were not a practice. New ideas require new behaviors and affordances to become a practice.

The cost of turning new ideas into practices and the uncertainty that the effort will pay off suppress adoption of potential innovations. The challenges of disseminating innovation throughout a community are explained in greater detail below, but in our framework, a maturation phase is needed first to overcome these barriers and turn ideas into mature

---

that he was involved in, to which I responded, "You're my advisor and you never told me this conference existed. What do you think?"

[41] I'm pretty sure that was a Jung metaphor as well… or maybe Proust.

practices. Where ideation involves microscale dynamics between individuals and their cells, maturation involves macroscale dynamics between cells and *cell clusters*.

Where an innovation cell is a densely connected set of individuals exploring a problem, cell clusters, as illustrated in Figure 3, are subgroups of cells within a broader community that collaborate on maturing a new innovation from knowledge into practice. Cells within a cluster might share domain interests, areas of expertise, and interpersonal connections (thick blue lines), but they must possess the same kind of complimentary diversity as individuals within a cell. Research on large organizations shows that when subgroups are semi-isolated from the broader community it prevents cognitive homogenization and herding around satisficing solutions, maximizing "equilibrium performance" of collective intelligence. Rather than rely on the paradoxical serendipity of moderate connectivity, matchmakers dynamically titrate connections within a cluster based on the incubation cycles of the individual cells. This dynamic moderation leads to "the highest equilibrium performance" of collective intelligence.

Where the received wisdom of density=serendipity is inefficient but functional when the spatial constraints of an office slow information flow, in distributed work it can be as crippling as no connections at all. The role of the matchmaker is to maintain optimal allostatic balance by removing and creating connections at just the right time. While it might seem challenging for matchmakers to intervene on a complex network, there are many examples of success. In neuroscience, groups are developing control strategies to target seizure suppression in the dynamic brain networks of epilepsy patients[42]. Back in the domain of social networks, analysis of temporal dynamics can predict social behavior with high precision or predict students' effort and performance from the social dynamics in a class. And research on innovation itself shows that matchmakers can actively induce intermittent interactions between cells that lift the novelty and impact of the best ideas. Informed by the microstructure within cells, matchmakers can strategically regulate cross-cell and cross-individual interaction. Titrating interactions maximizes the trade-off between independent thought and creative collaboration, preserving the average quality of work while lifting adoption of the best solutions.

Figure 3 traces the process of maturation. At Step 1, all of the cells are engaged in the early stages of ideation. There are various pre-existing social connections between individuals in different cells (dotted blue lines), but the matchmaker has paused direct communication. At Step 2, Cell A has developed a potential innovation, and maturation begins as the matchmaker opens up lines of communication in Step 3. Remember, at this point Cell A has invented broccoli. They have an exciting idea which is not yet a practice, and as we noted above, people resist investing in raw ideas. So, not only is this investment the cluster's explicit responsibility, but the matchmaker leverages social engineering to stack the deck in its favor. The specific connections activated by the matchmaker are not chosen at random, but rather, selected to maximize uptake by leveraging either high-trust, past working relationships (blue lines) or strong, novel personality matches (green lines). The goal of the matchmaker is to connect people that "get"

[42] Do you really think your company's social network is more complex than the brain?

each other, driving engagement with the raw idea. Even as the connections between Cell A and its neighbors are activated, the connections between Cell B and Cell C remain paused, allowing them to independently incubate the incoming discovery.

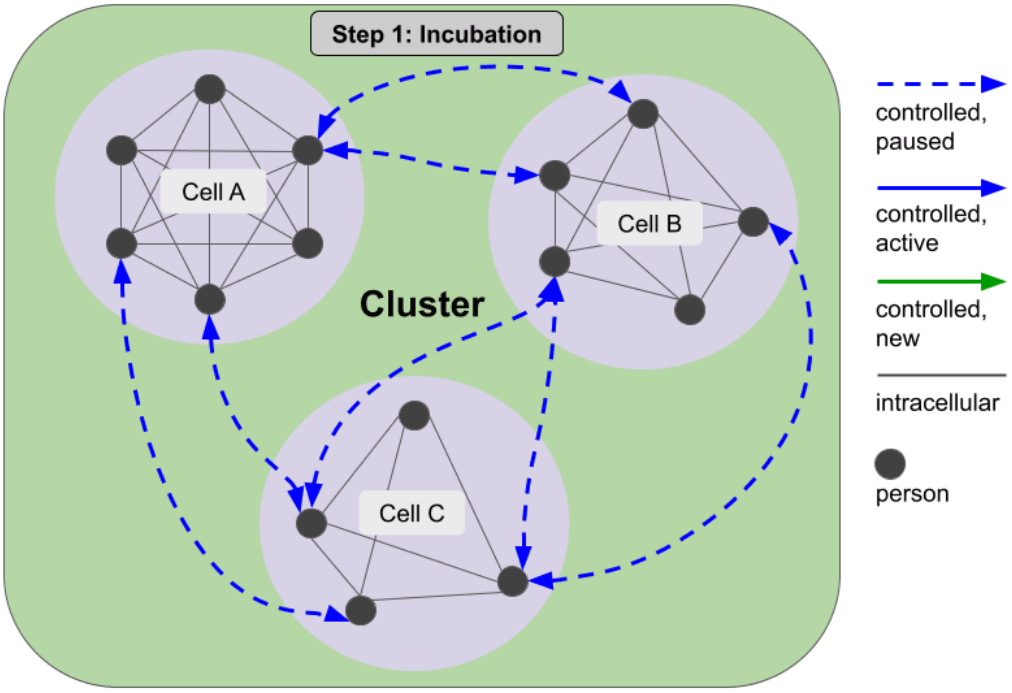

Fig. 3: The macrostructure of maturation.

In Step 4, Cells B and C have incubated their own variations of the original discovery. The diversity within the cluster allows it to identify new applications and hidden challenges associated with the new idea. At this point, the matchmaker opens all channels of communications between cells (Step 5), allowing the entire cluster to find an equilibrium in a set of behaviors and affordances which generalize across every cell. By Step 6, the cluster has matured Cell A's original idea into a robust practice ready for dissemination throughout the community.

In other words, Cell A invented broccoli. Cell B, noting the nutritional value when raw, developed a recipe for juicing it. Independently, Cell C came up with a reasonably tasty recipe for sauteed broccoli. Cell A noted that they don't have a juicer and that broccoli juice sounds pretty disgusting anyhow. Cell B argued that juicing it is the best way to preserve its nutrients, and so Cell C began exploring the optimal sauteeing time to maximize flavor while maintaining nutrition. Cells A and B, both with saute pans of their own, begin testing salt and garlic in their recipes. In the end, the cluster arrived at a tasty and healthy recipe requiring only tools and ingredients available to everyone. The idea of broccoli is now a practice of sauteed broccoli with garlic.

With a mature practice, a cluster can now disseminate a true innovation throughout the community. Just as with ideation and maturation, optimal dissemination relies on moderate

levels of connectivity and dynamic allostasis. This emerges, if rarely, even in the density=serendipity model. For example, the new social connections that result from individuals moving from high-impact smaller teams to larger scaling teams account for much of the global gain in innovation. Our framework turns this passive serendipity into intentional design[43]. The matchmaker actively engineers the macroscale dynamics by both remixing cell membership and modulating connections. One of their goals is to create "non-redundant connections" between cells, which play a crucial role in increasing adoption of innovations.

Successful dissemination is represented in Figure 4, Panel A. Initially, the matchmaker (lurking disreputably in its shady crimson triangle) has paused all of the controlled communication channels between clusters[44]. This allows individual cells and clusters to independently cycle through phases of incubation and maturation. In Steps 1–4, the central cluster develops a new idea, reaches equilibrium, and is ready to disseminate an innovation to the community. In Step 5, the matchmaker turns on communication only between the central cluster and its closest connections. With all of the other connections still paused, the central cluster can leverage trust and established relationships to reduce adoption uncertainty, and its three cells can leverage their non-redundant relationships to further drive adoption. As new cells adopt the innovation, the matchmaker strategically opens additional communications channels, creating a synchronized wave of transmission throughout the network (Step 8), which drives it towards equilibrium (Step 9).

Without the maturation phase or the intervention of the matchmaker, networks suppress innovation through both winner-take-all effects and high adoption thresholds. In Panels B and C, we see the same network of innovation cells attempting to develop and disseminate new ideas within a standard density=serendipity model. In Panel B, the initial cell has a novel idea which is immediately communicated to its neighbors without time for maturation. Without an established practice, it's neighbors in turn either ignore the idea or come up with their own variations on it. With no variation able to reach equilibrium, the existing practice (purple cells) eventually drives out the new ideas.

In Panel C, the original cell managed to develop a fully formed practice on its own. It immediately begins sharing this practice with others. Although they are more likely to adopt the fully formed innovation, this communication lacks the rich, non-redundant set of channels that the central cluster would have brought to dissemination. With most existing relationships still committed to the initial practice, the originating cell can't generate the synchronized wave of transmission needed to maximize adoption.

---

[43] Yes, we must replace the failed theory of the blind matchmaker with the sophistication of intentional design. I still, however, allow for the existence of a flying spaghetti inventor.

[44] As we've mentioned before, there are channels, like Facebook connections or personal friendships, which fall outside of a company's capacity to modulate.

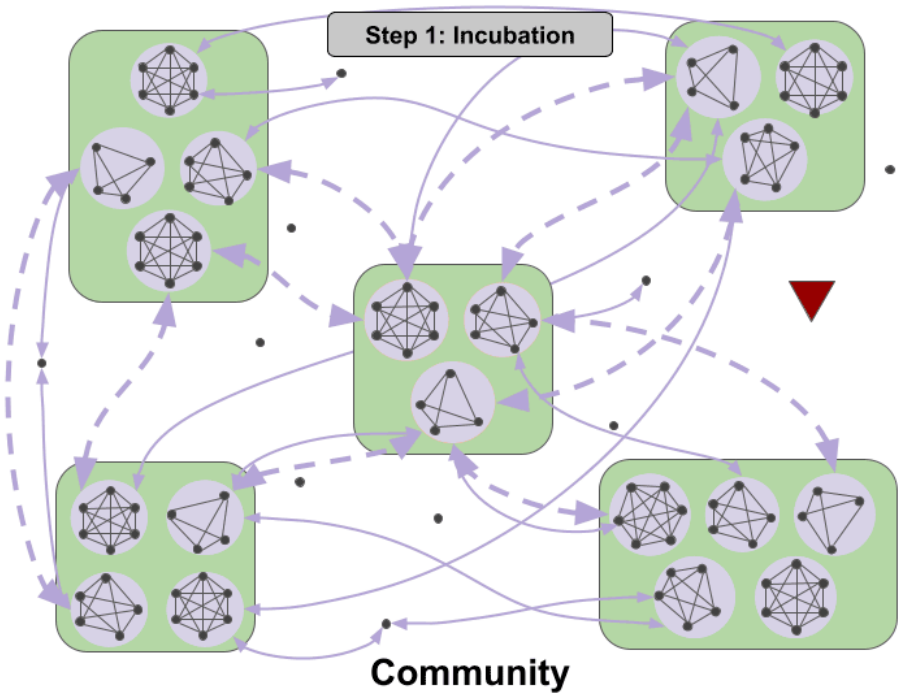
Step 1: Incubation
Community

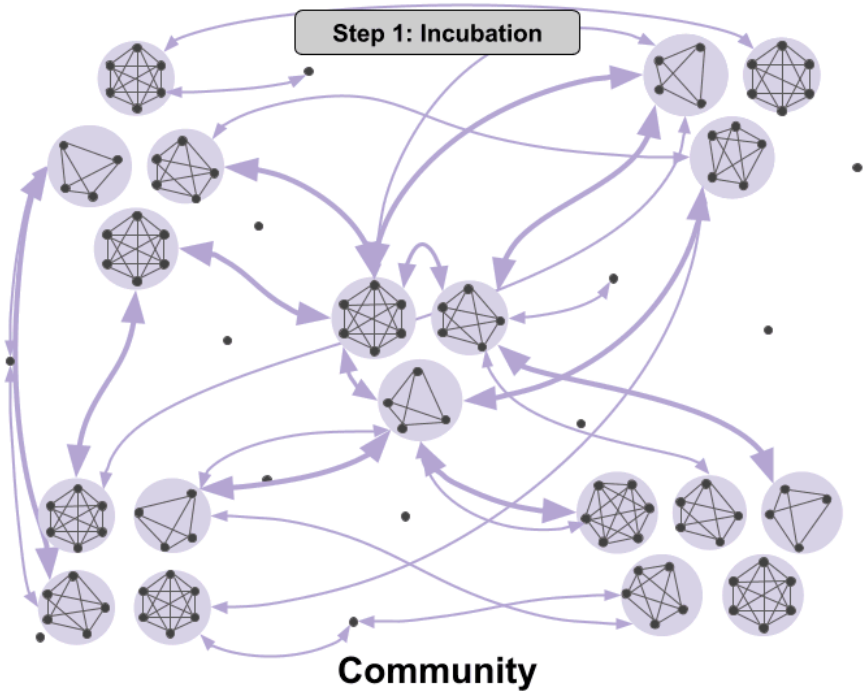
Step 1: Incubation
Community

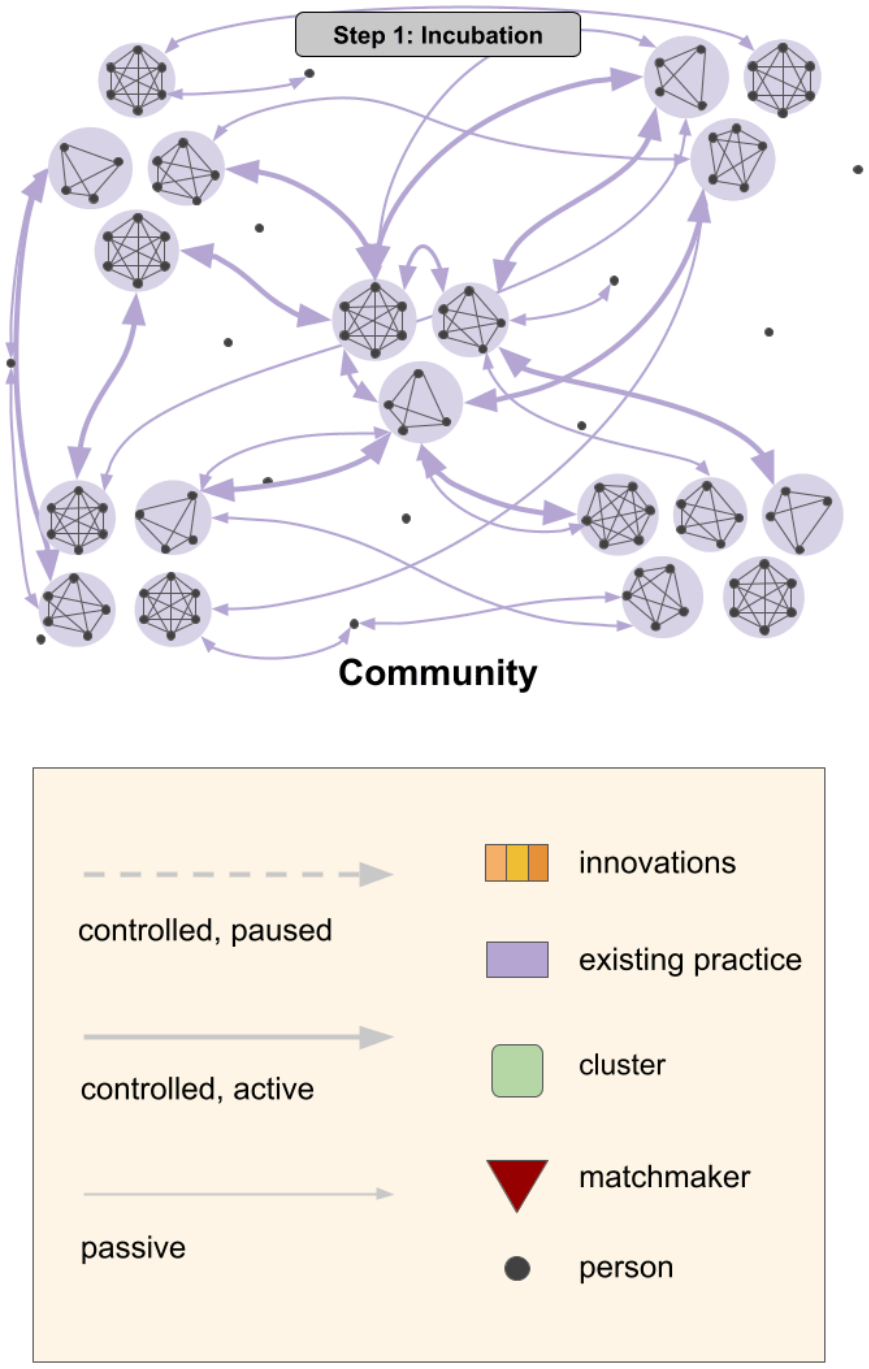

Fig. 4: Dissemination with and without maturation. (Panel A) Successful dissemination leverages maturation and matchmaker intervention. (Panel B) Without maturation, dissemination fails. (Panel C) Without clusters and synchronized communication, dissemination fails.

Socos's framework for innovation engineering introduces a number of crucial advances: tools that augment matchmaker interventions; the remixing of personalities and relationships

throughout a community; maturation of ideas into practices; clusters designed to leverage trust; and, matchmakers that peek in and orchestrate it all[45]. These innovations on innovation help to lift the rate of invention and dramatically increase the adoption of novel practices. All of this leaves one last question: how do you incentivize individuals and cells to maximize the collective intelligence of your organization?

Perhaps the most obvious starting point is a marketplace for innovation. Acemoglu and colleagues have posed innovation as a "resource allocation problem" where communities are confronted with aligning expertise and preferences of individuals and cells with the often hidden needs and complexities of unexplored problems. The idea can be thought of as a prediction market in which "dynamic pricing mechanisms...induce workers to self-select" into different problems that maximize their returns. If innovation were purely a skill-matching problem, we would expect market-based methods like this to maximize collective intelligence (defined here as the total number of mature innovations in full practice).

Unfortunately, open markets seem to quickly settle on satisficing solutions rather than substantive innovations. While this phenomenon is driven in part by densely connected networks and (paradoxically) easy access to information, it turns out that in innovation, market-based incentives "produce herding effects, reduce information available to the [community], and restrain collective intelligence." Each individual maximizes their returns, not by coming up with revolutionary ideas, but by letting the majority derisk their choices.

If you are feeling bad for Adam Smith, you might be happy to know that you can do much worse than an innovation market. A simple reward scheme in which individuals or cells receive fixed rewards whenever they develop a new idea substantially reduces collective intelligence compared to market incentives. But one incentive structure dramatically outperforms them all: "minority opinion". Here, rewards are given when individuals question the majority assumption and accurately predict an innovation few others saw. Wrong predictions and common predictions receive the same reward: nothing[46]. Minority opinion forces individuals to explore concept spaces far away from others, creating greater information value for the community. This incentive structure increased collective intelligence by 20% compared with market incentives and 40-50% compared with simple rewards.

In Socos's framework, minority opinion incentives work together with incubation and titration to drive exploration within and across cells and maximize the community's collective intelligence. Matchmakers inherit the same rewards as the cells they manage, encouraging them to diversify "bets" across a concept space. However, this entire scheme has the potential to produce unintended consequences. For example, scientific publishing has a well-known winner-take-all problem in which groups that are the first to publish (or often just the first to be recognized) receive all of the credit despite building on the insights of an entire field. This often induces

[45] Matchmakers are like the Illuminati, except that they didn't fake the moon landing.
[46] Innovation requires Boggle rules, but markets are The Family Feud… in so many ways.

scientists to be secretive about their research and inhibits sharing of data and methodology across cells. While minority opinion is not winner-take-all, only a minority will receive rewards for any given innovation. This presents a challenging problem of how to motivate the entire community when, by definition, the majority receives nothing and their unadopted discoveries are discarded. Instead, rewards should be given out not only to the cell and cluster producing an innovation, but also distributed to the entire foundation of insight that led up to the discovery. With asynchronous tools and concept mapping, we can understand how the inventing cell relied on past research and spread the reward back through the entire chain.

Think of it as “git-for-science”, where cells openly share hypotheses, methodologies, and discoveries to be forked and remixed by other cells. If a given repository ever produces an adopted innovation, credit and rewards are shared back through all of the contributors. In this way, unadopted discoveries contribute to the community’s collective intelligence and their originating cells are still incentivized.

Putting our framework all together, *innovation engineering* cycles through phases of *incubation*, *maturation,* and *dissemination. Incubation* launches the cycle with small, flat cells of diverse collaborators exploring concept spaces for new ideas. The *maturation* phase brings cells together into clusters to develop promising ideas into full practices. Finally, during *dissemination*, clusters coordinate communications to spread innovative practices throughout a community. Across all of these cycles, the *matchmaker* leverages tools like *personality models* and *dynamic concept maps* to *remix* collaborations and titrate interactions. They play the crucial role of maintaining the allostatic balance of trust and novelty within the microstructure of cells and across the macrostructure of the entire community. This *complementary diversity* combines with *minority opinion incentives* to maximize collective intelligence. Where the random chance of density=serendipity drowns out much of the very innovation it’s meant to facilitate, innovation engineering is an intentional practice that lifts the existing capacity of any organization.

You should stop relying on dumb luck and watercooler conversations as your source of innovation. Adopt our framework today or explore one of your own given everything you’ve just learned. As you do, you will discover one last major source of friction to innovation previously masked by the noise of density=serendipity. Throughout our framework, you can see the key role of complementary expertise, minority opinions, diverse concepts, and novel collaborations. The greatest innovation requires the greatest diversity at all scales. This brings us to the Diversity-Innovation Paradox.

Diversity-Innovation Paradox

I’ve already noted the large experimental research literature on the role of diversity in creative problem solving and collective intelligence. In fact, within science itself, simply adding new collaborators increases originality, innovation, and multidisciplinary impact, in part by restoring the allostatic tension between trust and novelty. Research like this extends these findings well

outside of lab-controlled experimental conditions, and this phenomenon extends beyond scientific innovation. Many years ago, I was on an advisory board for Credit Suisse, the Swiss banking giant, when we released a report showing that corporate boards with at least 1 woman had 3% greater return to shareholders, and that boards with 3 or more women returned an average of 5% more[47]. Research tracing the cultural diversity of thousands of London companies found that diverse management teams, in terms of both ethnicity and immigration status, produced more product innovations than their competitors. Even in the unsexy world of fishery ecosystems, diverse groups better understood "complex feedbacks and interdependencies" than homogenous groups. Just as with our matchmaker framework, however, the diverse inputs of fishery stakeholders needed a proper aggregation process–simply throwing every idea and stakeholder together in a room produced no advantage.

From global corporate boardrooms to London entrepreneurs, from fishery ecosystems to scientific innovation, increased diversity brings increased value creation. Still skeptical? What if we looked at all 1.2 million doctoral dissertations published in the last 50 years? This is exactly what was done in one of my favorite papers of 2020. Using text analysis and machine learning, the authors analyzed the career trajectories of every US doctoral recipient since 1977[48]. They found that "demographically underrepresented students innovate at higher rates". Unfortunately, this wasn't the paper's only finding. Those novel contributions from underrepresented scientists were less cited than their peers'. Despite contributing greater innovation, underrepresented scientists were less recognized and had worse career outcomes, robbing us all of their future impact. This phenomenon is known as the *diversity-innovation paradox* and it represents a major drag on global innovation.

Untrusted

While this paradox has deep roots in history, institutions, and personalities, it is fundamentally about trust[49]. As I've repeatedly argued, trust and diversity play fundamental and opposing roles in innovation and distributed cognition. Our well documented preference for trust over diversity drives organizations into the paradox. For example, cognitively diverse teams are more effective, and yet coworkers tend to sort themselves into cognitively homogeneous groups over

[47] Companies with all-male boards had an unfortunate tendency to preside over catastrophic collapse via malfeasance and bad acquisitions. Surprisingly, "Boys will be boys" is not sound corporate governance.
[48] This is where AI shines. Imagine what it would have taken for humans to read every piece of published research since 1977 and systematically rate them for novelty and impact. AI's "biases" reflect our own, producing a magical mirror that can occasionally speak hard truths.
[49] Untrusted nodes are a serious problem in distributed computing. "If you don't want everything to be taken down by one malfunction you need to defend against invalid inputs and being overwhelmed. Sometimes you also need to defend against actual attackers." The nearly trillion dollar computer security industry exists to protect distributed systems from untrusted interlopers through the simplest of solutions: trust no one. Whenever some developer forgets that rule, your credit card and social security numbers show up on a dark website in the "1,000,000 numbers for 1 bitcoin" bin.

time[50]. Even when an individual has a clear incentive to share information with a collaborator, they share less with individuals that are more culturally dissimilar. Left unchecked, these biases foment cultural barriers that suppress cognitive and broader diversity.

As cultural norms begin to favor homogeneity, individuals begin to conform to the norms of their ingroup "even when they understand that the norms…are arbitrary." Deeper still, their emotional responses shift to align with their perceived ingroup. Although there is significant variation in how individual brains process social signals around trust[51], a pattern emerges in which ingroup trust is associated with reward circuits while outgroup trust requires effortful top-down control. Additional stressors, such as time or financial pressures, increase the top-down effort needed for outgroup trust while having no impact on the reward circuits for ingroup trust. This creates another bias driving us towards homophily and undermining the prosocial behaviors necessary for success in diverse teams (and societies).

These barriers can be overcome, however, with engagement. Increased interpersonal experience shifts interpersonal trust from effortful medial prefrontal control to more reflexive (i.e., "automatized") brain circuits. Simple engagement with people of different races and genders reduces belief in stereotypes and other outgroup associations. It is only when trust and diversity are brought into balance that teams can avoid the Paradox and gain the full innovation potential of their community.

Without that balance, lack of trust dominates and robs us of capacity. For example, in many academic fields, including economics and the sciences, seminar audiences are systematically less trusting of women, asking more total questions and exhibiting more patronizing or hostile attitudes, even after controlling for differences in quality of work or field of research. In the world of entrepreneurship, investors of all genders strongly prefer pitches from attractive male entrepreneurs over female entrepreneurs, even when the pitches are otherwise identical. Analysis of meetings between venture capitalists (VCs) and female startup founders reveal that VCs undermine female founders with questions about losses while supporting male founders with questions about gains, and later negatively describe qualities in female founders that receive praise in male founders (e.g., "lacks experience" vs. "full of potential"). This difference in trust persists despite the finding that female founders return more than twice as much per dollar invested as male founders.

The potential anonymity of remote work suggests one possible solution to these known biases: we could all hide our differences within digital personas. Behind avatars and chat handles we

---

[50] It's interesting that even when "diversity" is not immediately visible, such as cognitive or intestinal flora diversity, cultural and psychological forces still tend to favor homophily.

[51] Recent research has suggested that most acts of explicit discrimination are carried out by only 5-20% of individuals.Much like Covid-19 superspreaders, perhaps a small number of people produce the biggest outbreaks. But also like Covid, the community as a whole suffers as these acts can become normalized or when this small population rises to powerful positions. How might 5-20% of professors drive large decreases in citations for innovative research by outgroup academics?

could all discard our identities and hide our differences. This has appeal for many, but rather than removing difference, all of the research above suggests that it would actually produce the opposite, exaggerating even subtle differences and worsening ingroup-outgroup effects. Trust grows with engagement, not anonymity.

As noted way back in "Measuring Remote Success", employees with more facetime with managers receive more recognition and more promotions. For most, though, working remotely reduces social network centrality, meaning that the average remote worker is less present in the minds of their coworkers[52]. This increases the need to manage one's impression amongst coworkers and leaders within the company. While everyone experiences this same increased need for "job-focused impression management", managers treat this behavior more negatively in employees of a different gender. Without the reflexive trust given to those similar to you, career management looks conniving and deceitful. This even plays out in elite levels of American politics–male senators are seen more positively the more they talk, while female senators…never know when to stop persisting.

"Distributed Innovation" explored how majority influence and herding effects reduce collective intelligence and innovation, and revealed the powerful impact of minority opinion incentives. This same effect is further exaggerated in situations demanding ingroup-outgroup trust. For example, when a deliberative process uses majority rule and few women are present, women experience more interruptions when speaking and their contributions are less trusted, reminiscent of the female academics described above[53]. When groups are required to respect minority opinion (e.g., unanimous rule), women are actually interrupted less than men, as the incentives of the system demand an effortful allocation of outgroup trust[54].

In "Controlling the Beyond-Control" I described how the attribution bias causes us to arbitrarily associate situational challenges to perceived failings of specific individuals. In the context of trust, it is clear that negative attributions become even more pernicious across group divides. Even location plays a role in exaggerating attribution bias as more distant collaborators receive more blame for failures, often hiding the real source of a failure and allowing it to persist. And for all our generic complaints of "Zoom fatigue", challenges with gaze and space in video meetings disproportionately affect those in outgroups. Cultural differences affect how we contextualize and interpret eye contact and other displays of emotion, influencing the way people interact. For example, East Asians tend to perceive direct gaze as evidence of anger or sadness compared

---

[52] A huge number of us have been working remotely for over a year, and many have asked me if we'll ever go back to the office. The shift may be slow at first, but eventually the most ambitious will want to move back to the center of the action, and when they do everyone else will be forced to follow.

[53] There has been a great deal of research on the complex dynamics of the classroom participation by female students. In many contexts, female students participate less than their male peers; these differences are further exaggerated in online and computer-supported learning.

[54] Wondering how to get started building norms around "minority opinion" or "unanimous rule"? Start with being completely transparent about the expectations: "...displaying the rules increased newcomer rule compliance by >8 percentage points and increased the participation rate of newcomers in discussions by 70% on average."

to Western Europeans. Even visual cues like the perceived distance or relative position of a face in an image affect perceptions of gender and emotion. And those perceptions in turn alter the neural circuitry of trust.

Even within America, factors such as differences in family size or self-expression can easily fuel attribution bias. For example, census data shows that Black families have 20% more people in a household than the average US family; Mexican-American families are 80% larger. This creates more competition of bandwidth and less private space for video meetings[55]. That little extra background noise or choppiness in video quality can easily lead to false assumptions about competence. All of this while that same video feed intrudes into a private home of those in outgroups, exposing "formerly safe, private spaces" to scrutiny and ingroup judgement.

Simple access to remote work is a form of privilege many don't fully appreciate. Even before the pandemic began, there were racial biases by income and industry in who was able to work remotely. Remote work also exacerbates broader socioeconomic and regional differences. Despite significant variation, remote "jobs typically pay more than jobs that cannot be done at home and account for 46% of all US wages." Of the highest earners in the US, 71% report being able to work from home, compared to only 41% from the bottom quintile. All of these differences in access further reduce the connections across socioeconomic, racial, and regional divides. In the end, you are more likely to encounter certain kinds of people working remotely, reinforcing in our minds that elite work is largely white-ish, male-ish, highly educated, and culturally homogeneous. The very tools we use to stay connected tend to exaggerate ingroup-outgroup effects and bolster the Diversity-Innovation Paradox.

Even the received wisdom of those with remote work experience reveals deep assumptions about the nature of trust and collaboration. "Set the meeting to 15 minutes by default, and only make it longer if absolutely necessary (the shorter the meeting, the more succinct you will have to be, and the less time there will be for pointless small talk and rambling)." A culture of 15-minute scripted meetings is a culture of homogeneity. Where is the trust in the assumption that others will waste your time? This idea might hold for rote cognitive labor where there's no need for chit chat, but in creative labor this mindset reinforces ingroup-outgroup divides and reduces innovation.

Distributed Fairness

I have spent the last several thousand words presenting evidence that diversity, combined with trust, lifts innovation and productivity. In fact, I've spent a good portion of the last 10 years arguing the business case for inclusion[56]. Scientists like to call this an *instrumental* argument,

---

[55] Less space per family member is particularly difficult for balancers that already struggle with intrusions between family and work life.

[56] I even have an upcoming book on this topic, *The Tax on Being Different*. Because you visited socos.org, our *terms of service* clearly state that you are now legally obligated to purchase copies for yourself and 12 friends. If, like me, you have no friends, you can choose to substitute a year-long

which is to say that there is a direct, tangible benefit to you. The instrumental argument for diversity and inclusion is real. It is more than valid. We all genuinely stand to benefit from overcoming the Diversity-Innovation Paradox. And yet, despite all of the force of this argument, it is not enough. Our choices must start with *fairness first*.

Fairness first means building a diverse team must come before all other considerations—it is non-negotiable[57]. It is not an argument about the potential business gains or the strategic value of balancing one objective against another. Fairness is about what is right. One might assume that the "rightness" of inclusion comes from its measurable value in innovation and beyond, but it is not enough to make a rational argument when so much of the research I've cited above reveals our deep and persistent irrationality in the face of difference. A recent study in the Proceedings of the National Academy of Science finds that instrumental arguments for diversity in universities appeal more strongly to White than Black families, and that universities that rely on instrumental arguments over fairness tend to have greater White–Black graduation disparities. When diversity and inclusion become mere considerations in a transaction, balanced against so many other pressing considerations, our long-term collective intelligence falls.

The constant demand to deliver new research, new products, new customers, new markets, places cognitive and emotional load on us all. Managing these daily stressors competes directly for the same neural resources needed to deploy outgroup trust, even as ingroup trust comes nearly automatically. Treating decisions about diversity as transactional might seem rational on its surface, one of many considerations you must balance for the good of your organization, but those same competing demands directly decrease our ability to see value in difference. This is why the Paradox persists.

If you believe, as I do, the copious research demonstrating that increasing the diversity of your organization will lead to increases in innovation and collective intelligence, you must then set that instrumental argument aside. Abandon the "business case for diversity". It is still true. It can be the argument that moves you to act, but it cannot be the act itself. Fairness first abstracts away from the transactional nature of the business case and simply states that fairness is a foundational principle to be supported independent of other considerations.

A few years ago, I developed an interest in how Hollywood balances creativity and industry, and so I began listening to *The Writers Panel*, a podcast about the "practice and business" of writing for television. The host interviews a panel of writers, often including a head writer, known as "showrunners". There were two fascinating commonalities in hiring practices of showrunners. First, nearly every one of them would say that there is an endless supply of professional writers

---

subscription to my upcoming Netflix comedy series, *I Wonder Who's Losing Their Job for Greenlighting This Pile of Shit?*

[57] If your immediate instinct in reading this is to reply, "But I don't want to hire a bunch of unqualified people to satisfy a diversity quota," then (1) of course they have to be qualified—how strange to assume they are not—and (2) diversity is a quality of a collection of individuals, not a single person, and teams that lack diversity are not qualified.

in LA who can deliver a script, but that what they actually needed was someone who could contribute an original idea the showrunner would never have had by themselves. This has been my hiring principle for years; I thought I was so clever and unusual, but I was just rediscovering what every good showrunner in Hollywood already knew[58].

Second, a number of successful showrunners would insist that they couldn't take the risk of hiring anyone other than friends and long-term colleagues. The theoretical value of diversity meant nothing to them against the daily pressures of delivering a high-quality show every week. One showrunner proudly stated that he would happily pay studio fines rather than diversify his writers room[59]. Another defensively argued the value of hiring only close friends while simultaneously complaining about how he had been marginalized for his own differences. Despite the universal agreement that new ideas are the currency of their industry and the causal relationship between those ideas and diverse rooms, the Diversity-Innovation Paradox persists, even in Hollywood. In one of the most competitive, high-stakes industries in the world, only those showrunners that practiced fairness first diversified their rooms.

None of this is meant to imply that fairness first comes at no cost. Building a culture that embraces the tension between trust and diversity takes time and effort. If a diverse writers room or innovation cell is non-negotiable, something else must give–usually, time. In the tech industry, the first qualified candidate who walks through the door is very likely to look and think like everyone else already in the room. Building a diverse team requires the time necessary to find a candidate that is not only individually qualified but also raises collective intelligence through complimentary diversity[60]. That inevitably requires more time and more effort on the part of recruiters, hiring managers, and the team as a whole. When a recruiter has 30 hiring managers every week scream at them for not filling open positions, you can easily understand how the instrumental value of diversity falls away as a consideration. Fairness first insists that a candidate is not qualified if they don't increase the complimentary diversity of the team.

The additional short-term costs of diversity don't end with recruiting. Even after a new member has joined a cell, outgroup trust is effortful, and that effort will inevitably reduce the short-term productivity of that cell. Expect to put additional work into maintaining psychological safety as you remix teams and recognize that this likely means short periods of decreased innovation as the room again finds its allostatic balance between trust and diversity. The long-term payoff for the initial effort will outweigh the lazy benefits of ingroup collaboration.

---

[58] I bet von Neumann was already doing this in his writers room 80 years ago.

[59] Writers rooms are a curious and somewhat unique practice to American television. In most other countries and other branches of entertainment, the economics doesn't support an entire team of writers dedicated to a single product. It is curious how this parallels many of the practices I've observed in innovation cells.

[60] This research, also from the lab of *Superminds* author Thomas Malone, shows that hiring the "smartest" or "most qualified" individuals without thinking about the whole group doesn't actually lift collective intelligence.

Many organizations don't allow the time needed to hire for complimentary diversity even as they talk about the importance of inclusion. It's one thing to understand intellectually that diverse teams and communities outperform in the long-run, but in a culture obsessed with short-term gain, the upfront costs of diversity can feel like losing. Fairness first frees us to see individuals as more than tools to complete this week's deliverables.

Of course, there are other strategies for increasing diversity. Many organizations offer incentives for "diversity" hires or include team diversity in managers' annual reviews. Research has shown that these strategies do have an impact, and yet the Paradox persists. It should tell us something about ourselves that we must pay people to make choices that are already in their best interest. I've been known to occasionally promise my kids ice cream and cartoons in exchange for a vaccine shot; doing the equivalent with grown ass adults just to build a qualified team is absurd.

Alternatively, some organizations fall back on the dirtiest word in HR: quota. On the surface, quotas might look like a form of fairness first–you must simply hire a certain number of a certain kind of person without influence from other considerations. And if people are using the q-word around your organization, there is probably a serious problem in need of a drastic solution[61]. But in practice, reliance on quotas isn't about fairness. Rather than making a commitment of time and effort to finding a candidate that helps build the best possible team, the appeal of quotas is that they are quick and easy. Just fill the diversity hire role. But diversity is not a quality of an individual hire, and focusing purely on individuals and not the composition of the team both ignores the importance of trust in the room and is an invitation to game the system[62]. Stop paying lip service to diversity. Stop making the easy arguments for diversity that ask little sacrifice and offer little gain.

Just as an individual must be qualified, teams must also be qualified. A team is not qualified if it is not diverse.

One Year Later

I started writing this piece in March of 2020. Now that it is finished, so is the lockdown, at least in some parts of the world, and many people will soon be returning to offices. What have we learned from our year of distributed work? And what will be its long term impact?

[61] Quotas, like democracy, are sometimes seen as the least worst option.
[62] For example, use the studio's diversity fund to hire a staff writer for your show. Then use her as a free writing assistant. Bonus, in the rare script that includes a character that looks vaguely like her, she can kick a line or two.

From Digital Natives to Distributed Natives

Companies can't rely on existing productivity measures to define success in a distributed world. Just as organizations have engaged with digital transformation, it's time for distributed transformation. I've laid out several starting points for this process:

- Expect metrics for baseline productivity to drop for most employees.
    - This drop is expected, not evidence of worker failure.
- Transition to metrics that measure adaptability rather than raw productivity.
    - Even better, leverage analytics that identify sources of adaptability.
    - Possible variables:
        - Type, duration, and number of employee interactions
        - Communication medium
        - Collective intelligence-related productivity metrics
            - Are individuals improving those numbers?
        - Interruption-related metrics and analytics
- Include "face-time" metrics to track interactions between managers and employees
    - Explicitly balance engagement across employees to prevent promotion based on the *availability heuristic*.

While I don't expect many large organizations to quickly transform into distributed work natives, many startups will. Much like digital transformation, those companies that engage early and make it an integrated part of the organization will experience outsized gains within their markets.

Optimizing Human Capacity

Don't assume that standard onsite methods for supporting employees are still viable. People and their circumstances are different; treat them differently.

- Balancers need structure, boundaries, and defined expectations.
    - Increase interaction with managers.
    - Leverage conspicuous monitoring
        - But don't overdo it; monitoring just one dimension is sufficient.
    - Establish unambiguous structure to their workday.
        - This includes setting hours of regular availability that mimics a normal workday without social pressures to deviate from it.
    - Create a unique workspace that provides a multimodal signal to your brain that defines "working".
    - Look for evidence of disengagement.
- Synergists need autonomy and self-regulation.
    - Don't layer one-size-fits-all expectations on top of their own self-management.

    - Give them more control over their own schedule and flexibility in their engagement.
    - Look for evidence of burnout.
- Recognize that different personalities (e.g., anxious or conscientious) will have best fits to different types of distributed work (e.g., routine and modular vs. complex and collaborative)
- Be honest with yourself and others about what you need.
    - Learn to recognize the clues that something is going wrong.
- Reduce distractions by batching emails and leveraging “office hours”.
- Invest in job training for idle workers.
    - Reap the human capacity benefits when demand for their labor returns.

One of the longer-term lessons from 2020 might be the importance of developing balancers into synergists. This lifts productivity, reduces mismatch, inefficiencies, and lagginess, and increases innovation. But this shift is effortful and might never include everyone. Those organizations that best support human capital development will see much stronger gains for distributed work.

Unfortunately, trends in automation and gig labor suggest that the greatest driver of adoption of remote work might be solely as a tool for cutting labor costs. For routine, modular work, the sort of work companies are already pushing into the gig economy, remote labor might become another means of lowering costs while also creating greater separation between “low-skill” workers and highly-valued “creative” employees.

Your Distributed Infrastructure

The transformation of companies from office-based to distributed organizations requires new tools and practices. It also requires a broad investment in the infrastructure that supports a distributed community.

- Minimize the use of synchronous communications (and make them count!)
    - Ensure that people don’t need to ask for information.
    - Use collaborative knowledge repositories (e.g., wikis, Notion, FAQs, natural language searchable databases, etc) to make all information readily available online.
        - Never trapped in one person's head or permissions
    - Go even further by offering knowledge repositories with dynamically explorable structures.
        - These promote cognitive flexibility and improve creative problem solving.
    - Enforce the use of communication that matches the urgency of the message.
        - Remember that chat (and sometimes email) is largely synchronous and can add to cognitive load and anxiety, increasing distraction and overload.
        - Emails must always include sufficient background, required outcomes, and an explicit due date.

    - Set a specific agenda with desired outcomes pre-identified for all meetings.
        - (But still allow room for natural peer-to-peer engagement to build trust and support incubation.)
    - *Please* train your employees in video setup and etiquette[63].
- Invest in asynchronous tools.
    - Don't force distributed teams to work as though they share an office.
    - Use asynchronous tools to develop *superminds*.
    - Minimize the use of tools that don't allow direct collaboration.
    - Never allow personal (hidden) notes or unresolved forks.
        - Align your understanding right there in the tool.
- Treat community broadband as a common good asset to your company.
    - Remember that variability in employee access is both beyond their control and a drag on your productivity.
- Supply your workforce with high-quality equipment for at-home work.
    - Develop a standard package of computers, cameras, lighting, microphones, wifi, and more.
    - Treat an investment in their home office as an investment in your infrastructure.
    - Guard against the corrosive effects of *attribution bias.*
        - Provide multiple channels of communication.
        - Allow employees to contribute in whatever medium is most effective.
        - Invest in public infrastructure.

I don't have much to offer as a prediction here. Infrastructure spending, including broadband, is already increasing, but it is unlikely to be either universal or robust. While it's possible we'll find ourselves in fewer meetings, haven't we been complaining about wasteful meetings for decades? Why would this be the magical moment for change? People often look at moments of disruption as though they are the fertile grounds of a profound shift in society, but this largely only comes true when changes are catastrophic. The industrialized world has (so far) largely weathered the storm of Covid-19 without sacrificing pizza deliveries, movie nights, or all-hands meetings[64]. Development of distributed infrastructure has experienced a big bump, but I'm not convinced it will lead to the sustained changes that many hope for.

Distributed Culture

Whether the world is in lockdown or not, never let chance, laziness, or bad habits define your work culture. Improving your tools can help with distributed transformation, but the real solution is a culture that is native to distributed work.

- Design teams based on core predictors of success.
    - Build teams around complementary diversity to maximize collective intelligence.

[63] I'm tired of being able to count peoples nose hairs on video calls.
[64] Some industries and communities have certainly been hard hit. But it seems like the net effect is that people that might have been underpaid to serve pancakes are now underpaid to deliver boxes.

    - Keep teams small and nimble to prevent homogenization.
    - Maintain flatter hierarchies within teams to promote more equitable contributions.
    - Promote prosocial perspective taking to increase psychological safety and productive risk-taking.
- Formalize goals, roles, and communication norms from the very beginning.
    - Review the specific norms I recommend in “rebuilding Culture”.
    - Use asynchronous tools to build superminds.
    - Promote norms that support individual differences.
- Make role-modeling transparent.
    - Share stories of the actions, sacrifices, and successes of both leaders and peers.
    - Remember, role-models that would normally influence the community are often hidden in distributed work.
    - Emphasize stories supporting minority opinion and psychological safety.

Humans have a vastly greater capacity for expressiveness and adaptability than machines[65]. I expect new cultures and norms to develop native to distributed work. In the sci-fi series *The Expanse*, humanity expands beyond Earth to both Mars and the Asteroid Belt. Over generations the “Belters” have developed a unique culture adapted to spending extended periods in space suits. Unable to see each other’s faces, they develop exaggerated arm movements and speech patterns to carry all of the subtle information previously carried in facial expressions. Mediocre cameras, bad lighting, spotty broadband, and tiny screens have made us all Belters. We will inevitably find new ways to convey the meaning of a grimace or eye roll. Soon, video chat will have its own emojis[66]. Or perhaps we all just become Richard Harris or Ian McCellan on camera and emote for the cheap seats[67].

As new distributed-native norms come to life, I suspect that they will follow the same rules I have identified of innovations: companies that use large, undifferentiated social networks will tend toward satisficing norms and ingroup-dominated cultures. The same tools we describe for engineering innovation can also be used to establish productive norms.

Innovation Engineering 101

Our year of remote work has revealed the unexpected truth that innovation is much more than density=serendipity. In fact, densely connected social networks and easy access to “answers” actually slow innovation. Whether teams are distributed or in-office, a dynamic allostasis between trust and diversity maximizes collective intelligence and innovation.

---

[65] For now.

[66] Some already do. I’ve given talks where “snaps”, “jaw drops”, and other vimojis fly across the screen as people respond to my pompous soapboxing. I want to build an AI that can do this automatically from peoples expressions–my sense of self-worth desperately needs to know they are laughing.

[67] I saw Ian McKellen in *Waiting for Godot* along with Patrick Stewart and Billy Crudup. They were all amazing, but McCellen’s face had total command of the audience. Every expression evoked a response. But if my wife and I could see every twitch from the balcony, he’d have probably looked like a crazy man in face-to-face conversation.

- Create a new role within your organization: innovation matchmaker.
    - Connect the right people and ideas, but only at the right time.
    - Build concept maps from asynchronous tools.
        - Use these to help prevent herding around easy solutions.
        - Connect ideas across innovation clusters.
        - Identify the right moment to make both social and informational connections.
        - Strategically regulate cross-team and cross-individual interaction.
- Remix groups of collaborators based on personalities to maintain the critical allostatic tension between trust and novelty.
- Make access to people and information intermittent to promote *incubation* and prevent satisficing solutions.
- Commit time for inter-team *maturation* of novel ideas into full innovations[68].
- Engineer a synchronized wave of communication to successfully *disseminate* innovations both inside and outside your organization.
- Leverage minority opinion incentives to drive collective intelligence.
    - Share incentives for “failed” innovations that laid a foundation for eventual success.
    - Avoid the herding and marginalizing effects of majority rule or prediction markets.

In many ways innovation over the last 20 years has stagnated. Distributed innovation can restimulate discovery and invention in our economy. I expect a substantial boost to innovation will also come from another source: secular reallocation. Over the course of the pandemic consumer spending and employment have both shifted dramatically away from certain industries towards others. Restaurants, bars, gyms, resorts, and even universities have seen dramatic reductions in their revenue and share of the labor pool. Companies like Amazon and Google received a great deal of the upside of that reallocation, gaining revenue and, in Amazon’s case, providing new sources of employment. There is an important pattern in this reallocation; the services provided by the “losers” have changed very little for 100 years, while the “winners” are some of the most efficient companies in the world. Spending $1000 on an iPhone will lead to vastly greater research and innovation throughout the economy than the same amount spent on meals at Cracker Barrel and visits to the gym. We might see a substantial lift in innovation over the next decade directly from this pandemic-driven shift in the economy.

Ending the Paradox

Diverse teams are consistently more productive when psychological safety is high, but when trust is low the contributions of diversity are systematically under-valued. Overcoming this paradox is crucial to distributed transformation as remote communication exaggerates the

[68] Don’t sell broccoli.

impact of ingroup-outgroup differences. Unfortunately, the roots of the paradox originate deep inside ourselves.

- Recognize that our very neural architecture supports ingroup formation and outgroup exclusion[69].
    - Online anonymity does not solve this problem.
        - It might even make it worse.
    - Promote interpersonal engagement to reduce outgroup bias.
    - Avoid majority rule (again) to prevent ingroup insularity.
    - Avoid massive, undifferentiated social networks for the same reason.
- Use the "business case for diversity" as a call-to-action but not as the action itself.
    - On its own, the instrumental argument fails to consistently improve diversity, inclusion, or equity.
- Practice *fairness first*: your teams are not qualified if they are not diverse.
    - Invest the time needed to discover candidates that are not just individually qualified but improve the complimentary diversity and collective intelligence of the team.

Neither Covid-19 nor the dramatic demonstrations for racial and gender equality around the world will "solve" the Paradox. I would love to be wrong, for this moment to be a catalyst for change, but no single event will truly bring a more inclusive society. We change with engagement and sacrifice. As the world returns to offices or adjusts hybrid work, a combination of cultural inertia and competing demands will sink the business case of inclusion just below the surface–visible but neglected. While nearly every public company has the goal of a diverse and inclusive workforce, few organizations are ready to place fairness first. The world will be transformed by those few that do.

A Perfectly Distributed Tomorrow

After a year of hand washing and false starts, some of us are finally returning to the office. But others are staying home, working from cafes, or "returning" to distributed officelets in formerly abandoned suburban malls. There will be much more to learn as we experiment with different versions of distributed and hybrid work. (For my money the future is hybrid quantum neural fuzzy distributed crypto work...in bed.) So, go forth and launch a Zoom meeting from your bathroom—just make sure you have good lighting and mute yourself liberally. Don't accept a workplace where the level of ISP competition in coworkers' neighborhoods determines the quality of the meeting. Do remember that lavishing praise on a female coworker you want to date doesn't count as a minority opinion incentive. And find comfort in the tension between trust and difference, incubation and maturation, synergists and balancers, synchronous and asynchronous, brownies and broccoli, near-term productivity and long-term human capacity.

---

[69] There are wide differences across individuals, but culture reinforces ingroup-outgroup formation even when only a small minority of individuals show explicit bias.

On second thought, don’t shit on camera.