THE ORGANIZATIONAL ANALYTICS E-BOOK

A GUIDE TO DATA DRIVEN ORGANIZATION DESIGN

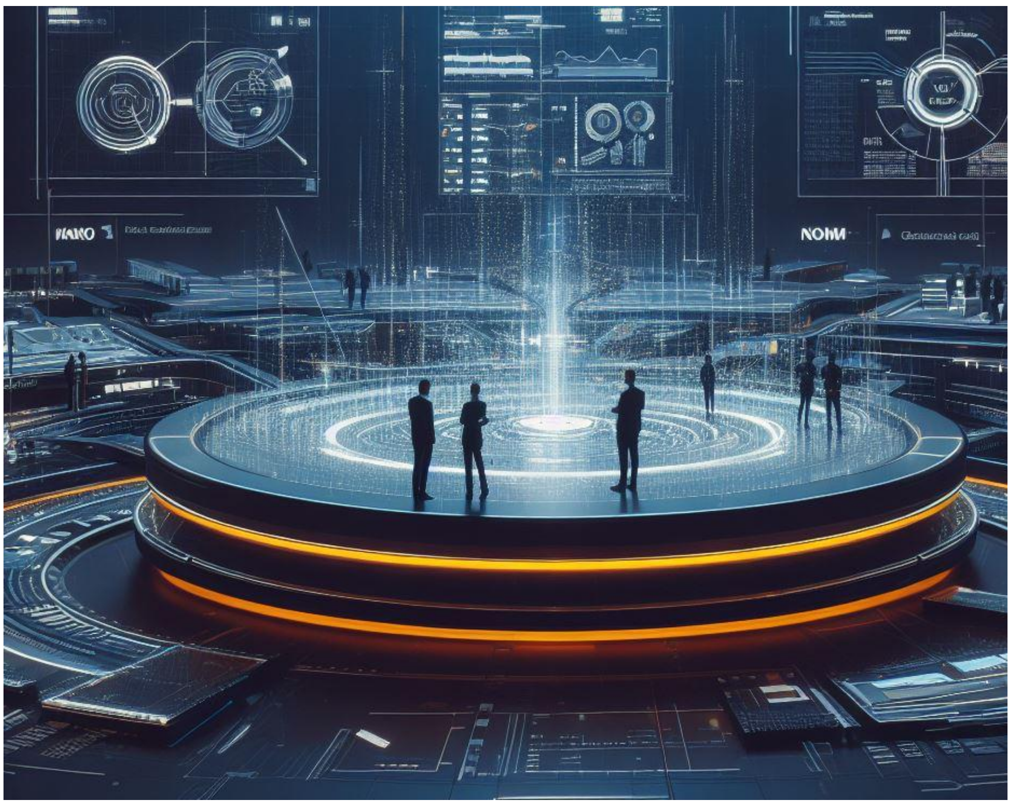

PHANISH PURANAM (INSEAD) JULIEN CLEMENT (STANFORD UNIVERSITY)

Version 2.0 (October 2023)

(Send feedback to Phanish.puranam@insead.edu or jclement@stanford.edu)

PREFACE

Knowing how to organize may be humanity's most significant and oldest accomplishment. But our lives today are also influenced by another, more recent, human achievement: algorithms. This book is about using the power of data and algorithms to design better organizations. We were inspired by our experiences teaching the **Org2.0** elective to **MBAs** at **INSEAD** and **Stanford**.

The ideas are drawn from our research in organization science, in which we have extensively relied on algorithms to analyze data from organizations and simulate organizations using computer models. The inputs of numerous industry experts who contributed to our classes as guest speakers and lecturers are also very gratefully acknowledged as they helped build the bridge from theory to practice.

This book is free. We'll be happy if you use it for teaching, research, consulting, or even entertainment. "How to organize" seems too important a problem today for us to put the little knowledge we have gathered on the topic behind a paywall, especially as we struggle to adapt to major shocks like the COVID-19 pandemic and the rapid developments in generative AI (such as ChatGPT and DALL-E2) aspire to make our organizations resilient to future shocks. We do request you cite the book if you use it.
The book is a living document— we hope to keep revising it based on your feedback. Do send us your thoughts: ***phanish.puranam@insead.edu / jclement@stanford.edu***

TABLE OF CONTENTS

CHAPTER 1

What This Book Is About

Organization design matters. Irrespective of size, sector, strategy, and goal, every organization must be designed in a way that helps its members pull together effectively to accomplish its objectives.

Organization design decisions come in all shapes and sizes:

- **how** (and whether) to form organizational units like teams, departments, and divisions;
- **how** to set up reporting and accountability;
- **what** incentives to use;
- **which** employees to hire, retain and promote;
- **where** to physically locate employees;
- **how** to best organize remote collaboration…

Notice that some of these topics are at the group level whereas others are primarily at the individual level. Both matter because an organization works well when individual efforts come together to reach group objectives. But in any given project we might put more emphasis on one level or the other. All of these are organization design decisions that will tremendously impact how people and groups in your organization perform today, tomorrow, and beyond.

ORGANIZATIONAL STRUCTURE & DEVELOPMENT PROBLEMS

Organization Design & Development Issues

Yet the way we currently make these decisions is, all too often, shockingly unsophisticated. That's why we wrote this book. Imagine having a headache and hearing the following from your doctor:

"I have what you need! Ancient wisdom says that herbs gathered under the harvest moon cures just your symptoms. Mr. X and Ms. Y, my other patients, have been using it regularly and they're really quite healthy for their age! Clinical trials? No there aren't any."

We wish this were a gross caricature. But this is roughly how many organization design recommendations are currently made, whether by managers within companies, advisers from consulting firms, or bureaucrats within governments. (Let's be honest; we probably did it a few times ourselves). Clichés such as "It's all about incentives" and "culture eats [pick your least favourite function] for [pick your favourite meal]" take the place of a science based approach to the problem. Trials to validate whether the recommended design can actually work for a particular organizational context are unheard of. Examples of "best practice" take the place of Evidence.

Copying industry best practice is a major driver of decisions in today's corporate world, and not just in the domain of organization design. This is a problem for two reasons.

- First, the evidence that practice truly worked well elsewhere is usually fig-leaf thin. ***Correlation is not causation***, and it is easy to be fooled by randomness into thinking that a particular structure or practice drives success when in fact it does not (or perhaps even harms performance!). Mr. X and Ms. Y's health may have nothing to do with the pills they are popping, or worse, perhaps they are healthy despite them.

- Second, ***all organizations are different***. Just because something worked well elsewhere does not mean it will work for your organization. Even what really keeps Mr. X healthy (or at least does no harm to him) may be poison for you. At best, you may be paying for useless pills. The biological variation that exists between people is almost nothing compared to the social variation that exists between organizations, so we'll let you imagine how often organizational practices imported from one company may really be poisonous to another.

Of course, you might say criticizing is easy. Relying on industry best practice has its problems but perhaps there is no better alternative… is there? Yes, there is! Recent developments in two areas have changed what is possible today in the field of organization design.

- **First**, **developments in algorithms and the availability of "Big Data"** have enabled us to observe much more of the internal workings of companies and analyze them with far greater sophistication than was possible before. Critically, the data and associated methodologies allow us insight at a micro-level into how individuals are affected by the organization designs they live in, as well as how their behavior aggregates up to produce the outcomes of their organizations.
- **Secondly, new theories of organization design** have emerged that move the action away from the "boxes and arrows" of the organization chart and instead help us think about design at the level where the data and indeed the innovations in organizing now reside: individuals, teams, online communities, and networks—all of which may exist within the legal boundaries of a large corporation or government department, or indeed cut across them. Together these developments enable an organization to analyze its own data to design itself.

An article and short video on the contrast between old and new approaches to organization design

EMBEDDED LINK

We use the term **"Org2.0"** to refer not only to radically new organization designs that are emerging today in response to developments like the post-millennial workforce, advancements in AI and the COVID19 pandemic, but also to the **suite of methodologies** and thinking tools that enable organization designers to make much more sophisticated decisions than they could in the past (with "Org1.0" if you will). We can use these new techniques to better perceive the organizational environment of a company, make predictions about specific events in specific parts of a company, and experiment quickly but rigorously with new practices to choose the ones that work best. These are **tools for organizations to design themselves**, rather than rely on designs copied from elsewhere. With these, we want to show you **how to design, *not what design to use*.**

New designs emerge all the time - **"*flat firms*", "*holacracies*", "*all remote*", *decentralized autonomous organizations (DAOs)*"** and **"*agile*"** are the ones in vogue as we write this- but they often wax and wane (see the last chapter for some of our thoughts on these). The shift towards ***data-driven design***, on the other hand, won't be going away. The goal of this e-book is to give you an intuitive understanding of the analytical tools which make this possible, provide a framework to link these tools together and show you why (and when) they can lead to better decisions with sizeable benefits. And all that in 7 very short chapters! - and quite a few embedded articles and video links if you're curious.

Reading this book won't turn you into an analytics expert, but it **should make you fluent enough to interact with and evaluate the work of analysts whose reports will support your decisions.** The ideal data science-based organization designer would be a bit of a Unicorn- in this widely shared image from IBM labs, you can see that it's hard to find people with all the required skills.

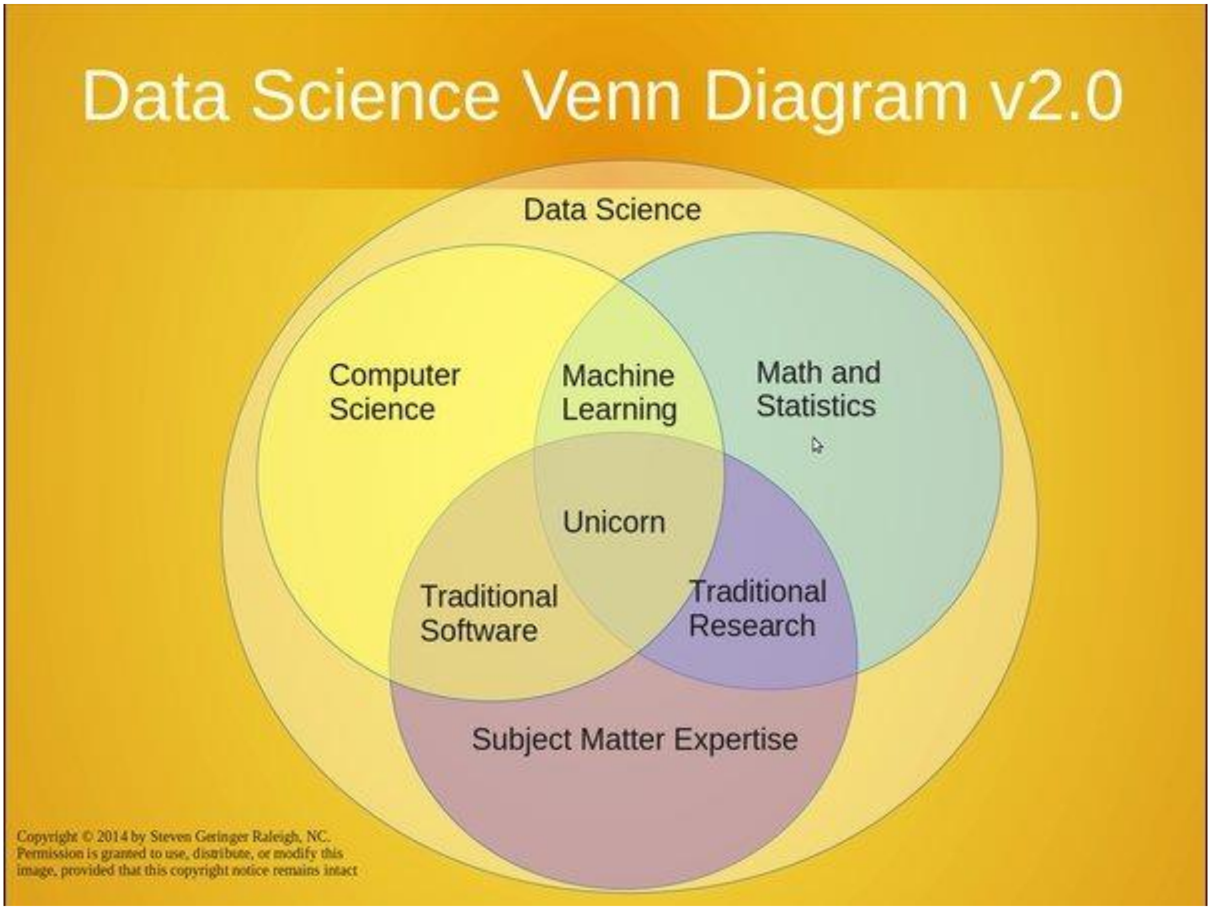

But we can build teams that collectively have these skills, and one of the basic things we know about teams is that they are able to work better together when the members know enough about what others know and how to use each other's expertise.[1]

Our book is written primarily for the **"*subject matter experts*"** in the picture above, and the managers whose teams hold these different skills. Accordingly, our approach will be to "black-box" the technical detail in the text: we'll focus on *what* analytical techniques can do for you rather than *how* they do it. The embedded video lectures and links provide extensive details on several analytical methods for curious readers.

[1] Moreland, Argote and Krishnan (1996);Reagans, Miron-Spektor, & Argote, (2016)

We have prototyped and refined this approach extensively with MBA students and senior executives we have taught at INSEAD and Stanford. It seems to work well… but in the spirit of this book, we hope you will experiment with it before deciding. For the same reason, this book contains no detailed case studies about exemplar companies that are using Org2.0 approaches in their organizational development work. We know first-hand many companies that do, and many of them have visited our classrooms, but using them as examples to argue our case would constitute nothing but the “best practice” fallacy all over again. So, read on and make up your own mind!

A short video providing an overview of the Org2.0 approach

EMBEDDED LINK

Reflection question: why do you think “people” related decision-making in organizations has been slower to rely on data and evidence, compared to decisions about customers, finance, and operations?

CHAPTER 2

The core problems of Organization Design

Overview

- It's useful to think of any group with a goal as an "organization". That includes teams, departments, divisions, and entire companies.
- Every organization, regardless of size, location, and industry faces the same universal problems- how to divide labor and then integrate effort.
- While the problems are universal, the solutions are not. Finding solutions that work for your context is the essence of organization design.

At the most basic level, Org2.0 is just a particular approach to solving the age-old problem of organization (re)design: how to create systems of individuals who may differ in what they know and personally want, that can still accomplish a goal together. That involves a few fundamental problems which are neither new nor necessarily unique to human organizations. Ants face the same challenges, as do computer networks, and squads of robots working together- though as far we know, none seem to rely on management consultants to solve them...

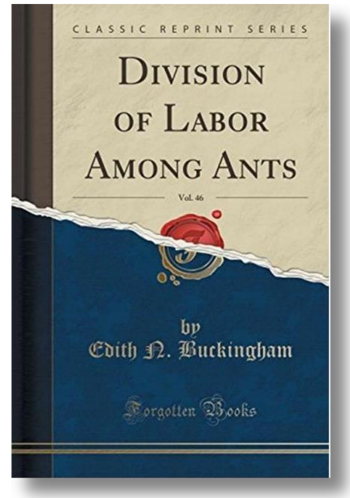

What are these fundamental problems? ***Division of labor*** and ***integration of effort***.[2] Roughly, these correspond to "breaking things down" and "putting them back together again": organizations divide their activities into different tasks performed by different people and groups (often because specialization yields efficiency), which then need to coordinate to produce a coherent output.

We can detail it down further. Designing the *division of labor* requires partitioning the overall goal of the organization into **sub-tasks (*task division*)** and assigning **sub-tasks to individuals and groups (*task allocation*)**. Designing *the integration of effort* also involves dealing with two related sub-problems: motivating individuals through **money, autonomy, recognition, and other means (*reward distribution*)** and making sure they are sufficiently informed about **what they need to do (*information flows*)**. And since no set of solutions to these problems is likely to be either perfect or permanent, we need some mechanisms for re-solving these problems when exceptions arise **(*exception management*)**.

Embedded Link: The universal problems of organizing illustrated

EMBEDDED LINK

[2] Mintzberg, (1979); Burton and Obel (1984); Puranam, Alexy and Reitzig (2014)

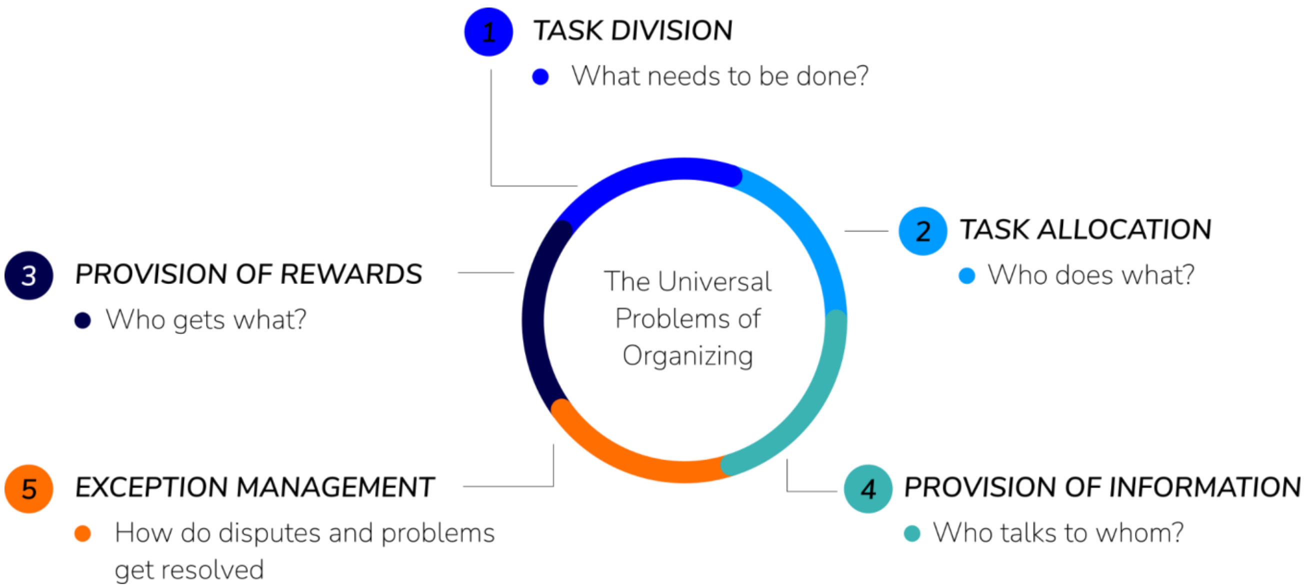

The Universal Problems of Organizing

When we (re)design an organization (and we stress "re-" because most of the time, we must change existing designs) we are always trying to answer the same basic questions. Given the goals of the organization we are trying to design.

- **what needs to be done** (task division)?
- **who does what** (task allocation)?
- **who gets what to motivate effort** (reward distribution)?
- **who knows what in order to ensure competent and coordinated effort** (information provision)?
- **what to do when exceptions arise and new solutions are needed?** (exception handling).

To summarize: **designing an organization means searching for solutions to these universal problems.** Improving the effectiveness of an organization (a.k.a. redesign) is equivalent to finding better solutions to the universal problems of organizing that confront the particular organization. This will be true in every existing organization you know of, whether it is an "old economy" manufacturing company, a cool new technology start-up, a software development team, a football team or a string quartet. It will equally be true when you are trying to figure out whether the entire company should be reorganized, and when you are simply interested in how to get a particular team to work more effectively. That is the beauty of thinking in terms of universal problems! The problems are always the same, and the differences will be in the solutions deployed by different oganizations: different solutions work for different organizations, which is why you should always be very careful in asking yourself whether what other companies are doing will also work for you.

Reflection question: Think of any group that you have recently been a part of, at work or home, that had a specific goal or task. Can you explain what solutions this group had to each of the core problems of organizing? How effective were these solutions?

CHAPTER 3

Organization Design Analytics

Overview

Analytics-driven approaches to organization design can be conducted at three levels of sophistication:

- *Perception*- describe what's happening now; this is where Big Data combined with traditional statistics have been very useful
- *Prediction*- forecast what's going to happen before it does and – this is the domain of machine learning/AI applied to Big Data
- *Prototyping*- decide which intervention has the best chance of producing the results you want - this is the world of A/B testing.

You don't always have to go for the most sophisticated approach. For some problems, it might make better business sense to just rely on perception or prediction. You have to consider the *value of* the *information* generated through prototyping.

What does an analytics and data-driven approach to organization design – finding solutions to the universal problems of organizing that are applicable in your specific context- look like in practice? The analytics-driven design comes in three levels of sophistication: perception, prediction, and prototyping.

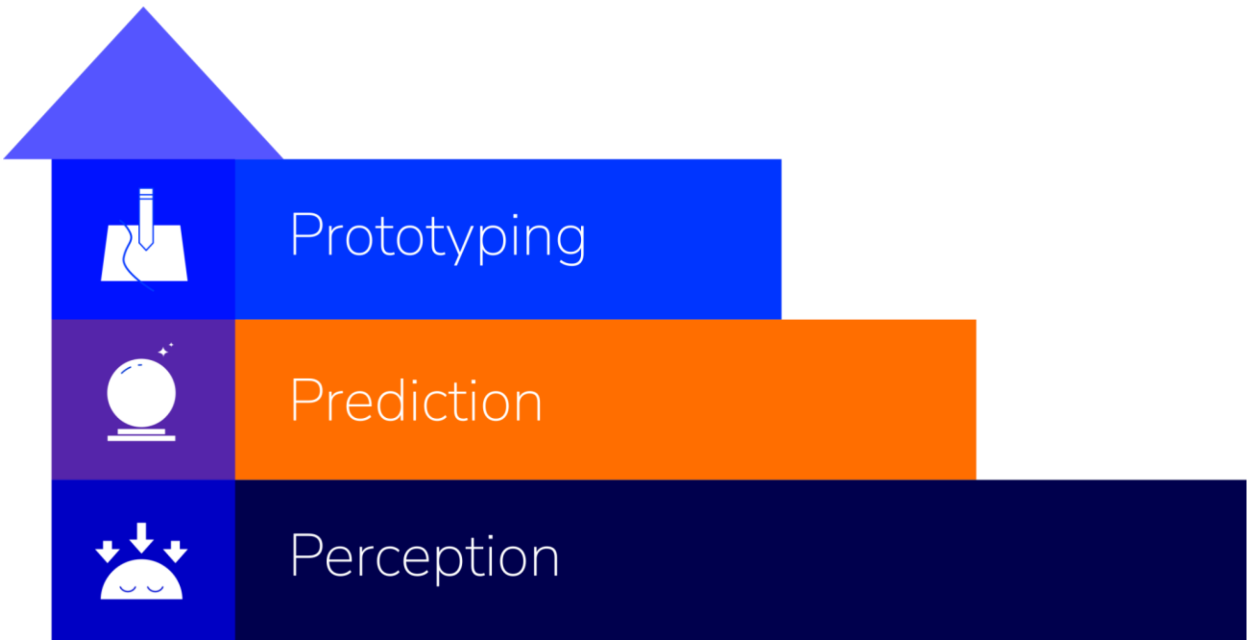

We'll use a typical organization design challenge as our running example: the problem of "silos" (i.e., absence of coordination) between different units within an organization. Silos are often blamed for delays, missed opportunities, cost inefficiencies, and lack of innovation. How would an analytics-based approach tackle this most fundamental of organization design challenges?

The first step towards better design is ***perception.*** Perception is about the present. It involves forming an accurate picture of how the organization is working currently. A major requirement in our case is to gather data that allow us to answer some basic questions, such as:

- How frequent are the instances of breakdown in collaboration?
- What is the (estimated) economic impact of these problems?
- Where do they seem to occur most frequently?

Many apparently "burning issue" design challenges may actually turn out not to be so important after all; as we will show later, "silos" are not per se a problem. In fact, it's nearly impossible to organize a large company without some silos. The trick is to have the silos where they should be, protecting from distraction the things that need such protection, but allowing interactions between people who need to interact to get things done. To figure this out, we need data. Archival records, surveys, and even interviews can help us see if there really is a systematic problem. All too often, we are just listening to one irate (and vocal) senior manager who feels their project did not get the attention it deserved from colleagues in other departments. When we have a large enough sample of data, we can run perception analytics (summary statistics, cross-tabulations, and hypothesis tests for example) to assess whether the data are showing a pattern or just noise.

A video lecture on the basis statistical tools for Perception

EMBEDDED LINK

Good data also enables us to perceive problems before they become widely visible (or because they are being masked by politics). Are certain interfaces between departments particularly prone to friction? Are the "missed opportunities" created by silos imposing a large opportunity cost on the organization? Sometimes, data is useful to simply debunk myths around any of these types of issues too. Consider "Those folks in marketing live in their own world!" vs the data showing that most of the breakdown in handoffs are occurring between design and manufacturing. Data can also help to generate insights and test hunches. For instance, your data might show (or you might be able to test your hunch) that silos are particularly damaging when they coincide with geographic or time-zone boundaries. Put simply, the analytics of perception is first and foremost about making sure we are perceiving the right problem- it is the equivalent of a thermometer or a blood test to detect malaise, as opposed to judging from looking at the patient or holding a thumb up to their forehead, or just asking them how they feel.

The next level of sophistication involves ***prediction***. Here, we move from understanding what has led the organization to work in a certain way until now to predict what will happen to it in the future. This is like a blood test or an X-ray – it helps predict what happens next, not just describe what is happening now. In doing so, we can also move from simple aggregate patterns ("Legal department is routinely holding up approval on bids our engineering team is making, and we are losing business") to forecasts about specific events ("*Tom*'s bid that he needs to get to the client in 15 days is going to be held up at legal cell with a 69% risk of failing to meet their deadline!"). Making accurate predictions can be very powerful because **ultimately all decisions rely on predictions: accurate predictions lead to better decisions.** Consider this: every time you make a decision, you choose between alternatives, based on your beliefs about which will produce a better outcome for you. That's a prediction problem, and in many situations machine learning algorithms (the relevant branch of AI) can help build predictions (for instance, by looking at historical data about projects and approval timelines).

A video on the basic machine learning tools for predictive analytics

EMBEDDED LINK

Though very useful, machine learning isn't the most sophisticated technique you can use for organization design. That's because prediction isn't the same as explanation: just because we can predict that something is likely to happen doesn't mean we know *why* it may happen. For instance, we might be able to predict that contracts written by a certain sales manager are likely to face more delays at the legal unit than those written by other managers, but that doesn't mean we know whether this is because these contracts are poorly prepared or because the legal unit lacks expertise about these types of contracts. Org2.0 thinking, therefore, suggests that we use predictive analytics to produce ***indicators*** that something is worth investigating ("let's keep an eye on the contracts that are forecast to go slow and put some energy into chasing those") rather than directly as ***interventions*** ("these contracts are likely to take forever to clear, let's replace the managers!"). Indicators could be variables that are good predictors of other variables, and the predictions themselves can be useful as an early warning system that something is likely to happen. An intervention on the other hand is a manipulation of a variable that causes a change in another variable. Indicators can give us good hunches about interventions but should not be confused with interventions.

There are three other reasons why we advise caution with machine learning. **First**, predictive accuracy often comes at the price of an intuitive explanation. More complex (and less explainable) algorithms may make better predictions in many situations but leaving our decisions entirely to "black-boxed" algorithms can be highly problematic when there are significant human consequences to the decisions taken as a result. Algorithms optimize what we tell them to optimize (for instance, "find strong correlates of who is likely to be a high performer in the data") but not any of the criteria we usually leave implicit in our decisions ("make sure we give equal chances to employees from different backgrounds"). They learn from historical data that may reflect past biases in decision-making, and if used blindly may perpetuate such biases.

Second, human beings are hardwired to try to understand *why* things happen. People are much more likely to accept and support decisions when they understand the reasoning behind them. Telling them that "the random forest algorithm says this is likely" won't cut it. This is why predictions are a good start for decision-making processes, but this is not an end in itself.

Third, just having a lot of data does not guarantee that machine learning can work for you. There are a number of potential challenges with data, that even in large volume, may make it useless. In essence, all machine learning techniques work on the premise that data gathered today are representative of how the underlying processes generating the data will unfold tomorrow. That may be false for a variety of reasons. In general, **the stronger the resemblance between the situations you are gathering data from and applying it to**, the better.

[A short article on why lots of data don't always translate into success with machine learning]

EMBEDDED LINK

The last design tool, ***prototyping***, is the only method we know that helps understand what *causes* the relationships we observe in data. Here, we don't rely on past data but rather create new data by running "experiments" within an organization. Experiments involve randomly assigning some units (i.e. people, teams, projects, bids) to a treatment condition (the new policy we're thinking of implementing) and others to the control group (where things stay as they were without the new policy).

Randomization is crucial! Imagine that you implement (without randomizing) a new training policy and find that employees who applied for the training and attended it saw a rise in their performance evaluations. You have no way of knowing whether this is because your training was effective or because the people who applied for it are just generally more motivated high performers whose evaluations were going to rise anyway. You might object that in this case, you might just make the training mandatory for all employees and see whether everyone's performance went up afterward. But then, how do you know that the increase in performance wasn't due to something else which affected all employees, perhaps because demand for your products went up this year for unrelated reasons? Randomization allows you to avoid all of these problems by creating ***counterfactuals***, that is to understand what **would have happened** without the intervention. This is possible because randomized treatment and control groups are ***statistical twins***: they are similar enough to be treated as identical, so the control group can serve as the counterfactual. **We cannot establish causation without counterfactuals, and we cannot establish perfect counterfactuals without randomization (unless we have time machines).**

Randomization can help you understand whether any new procedure (not just training!) is effective. For instance, we could imagine solving silo problems by setting up a new procedure to process contracts by the legal unit. How do we know it will work? We pick a random subset of contracts to apply the new process to, let the rest go ahead as usual, and compare outcomes later.

That's it! For instance, to solve the silo problem, should we invest in team building? Change the staffing of boundary-spanning roles? Create a fast track for certain projects that have to cut across departmental boundaries? Rather than implementing one of these based on our hunches, we could prototype them and only roll them out widely if the evidence looks promising.

There is one other advantage that prototyping enjoys. Prediction and perception are both hungry for data: they require you to collect data not only on the outcome you're interested in but also on a whole set of variables that correlate with it (i.e., the X's that associate with the Y's).

But prototyping need not be. If you randomized the treatment properly, all you need is to measure differences in outcomes (Y's) between the treatment and control group. This avoids a whole lot of concerns about data privacy that often come up when using the Xs in prediction or perception.

Video lecture on why an experiment with randomization is the gold standard

EMBEDDED LINK

Of course, experiments with randomization are costly. This means that they should be used only when the cost of "getting it wrong" is large (i.e., when the potential benefits of your new design are large and should not be missed out, and the potential costs are large too if it turns out to be a failure). If either the benefit of succeeding is large, but the costs of failure are small, or vice versa you don't need an experiment.

	Cost of Failure: Small	Cost of Failure: Big
Benefits of Success: Small	Ignore	Abandon
Benefits of Success: Big	Implement right away	Prototype (e.g. Experiment)

There are a host of other tricks we can use to create reasonable counterfactuals when we can't resort to randomization in the organization. Scenario discussions and computer simulations, for instance, can help us understand which intervention will be effective to solve our silo problem. One might gamify the task that people do (i.e., create a toy version of the kind of project they usually work on) and run a hackathon over a few hours or days, having different groups try to solve the problem after randomizing them into different interventions that can help solve the design problem we are tackling. For instance, are agile teams good for us? We could randomly assign some teams to agile structures while others stay in the current design and all teams compete in a 24-hour hackathon to complete a simplified version of a typical problem.

A short article on Gamified Randomized Control Trials

EMBEDDED LINK

Prototyping with randomization involves experiments *in vivo* (literally "within the organism"). An alternative is an experimentation "*in silico*" through computer simulations. Simulations allow organizations to create representations of how they work in the present and explore a variety of "what-if" scenarios to understand how to ensure better coordination between their members, how to convince employees to adopt new management practices, and many other important processes involving interactions between many agents.

An analogy may be useful: a traffic designer or a bridge designer is unlikely to use randomized experiments to prototype a new system of traffic lights or construct a new bridge (we hope) but will instead rely on running the experiments within computer models of the system. If we had a good computer model of how bids get approved, for instance, we could see what the impact would be of changing a particular step in the process. This is no different from trying to figure out what might happen to your projected Net Present Value (NPV) for an investment if you change the expected growth rate of revenues or the cost of capital.

Video lecture on alternatives to randomization for prototyping

EMBEDDED LINK

Reflection question: when do you think a Prediction approach (e.g., "Who is likely to quit next quarter") is superior to a Perception approach (e.g. "What factors correlate with people quitting?")? When is prototyping superior to both?

Bonus Reflection question: the evidence suggests that many if not most organizational change initiatives fail. Yet prototyping a change before full-scale implementation seems rare (how many companies do you know that use clinical trials to decide on their next re-design?) Why?

CHAPTER 4

The first lever of organization design

Overview

- Structure refers to stable patterns of interaction. They can be laid down in rules or emerge organically
- Interactions come in two types: interdependence – A needs B's inputs to get her job done; and influence-A's beliefs and actions are shaped by B's.
- The most basic principle of structure in organization design is to make sure that "influence matches interdependence".
- To fix a design problem through structure, therefore, requires identifying and correcting mismatches between interdependence and influence.

When they think about the levers of organizational (re)-design, most people think immediately about structure: the "boxes and arrows" of the organizational chart. This view is doubly narrow. The structure is not the only lever of organization design, and structure is not only about the boxes and arrows. The research is quite conclusive that individual behaviour in an organization is jointly shaped by three levers of design: structuring, sorting, and sensemaking.[3] Therefore, we suggest that organization design analytics can potentially operate through each of these levers individually or in combination, though every approach is not equally feasible in all situations.

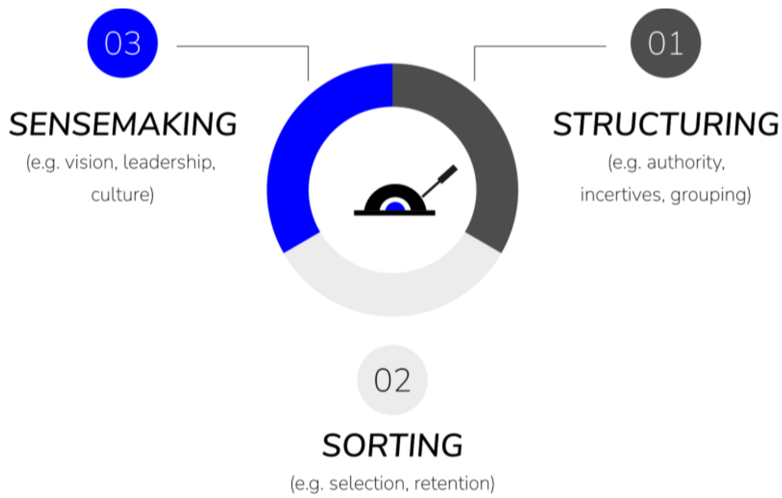

THREE LEVERS OF ORGANIZATION DESIGN

"Structure" simply means "stable pattern of interactions between people". There is a lot that goes into that! They occur at all levels and layers. They involve peer-to-peer as well as superior-subordinate relations. Some interactions are mandated by somebody with authority—they are part of the "formal" structure—and others emerge on their own—they represent the "informal" structure. Interactions between people are the lifeblood of an organization.

Organizational charts only show interactions that are part of the formal structure (the figure below is the first documented instance of an organization chart, used by the New

[3] For instance, see the recent review by Chapman and O'Reilly (2016). The McKinsey "7S framework" is also broadly consistent with these ideas.

York and Erie Railroad from the 19th century). For instance, a product division in an organization chart is shorthand for **“all employees in any function who contribute to making and selling this product are expected to prioritize their interactions with each other (over interactions with those in other divisions) and be subject to a common source of authority (the director of that division)”.**

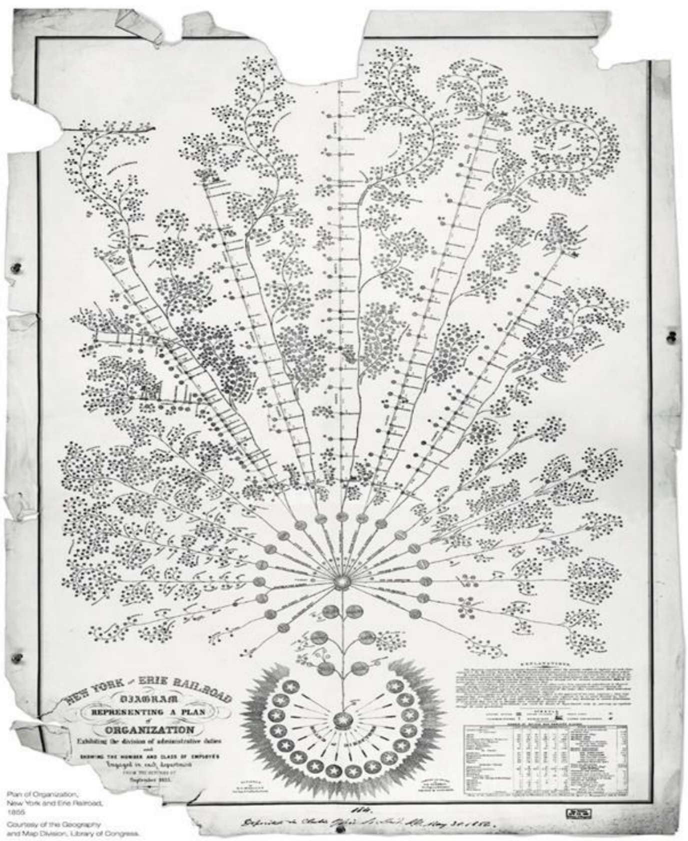

Plan of Organization, New York and Erie Railroad, 1855

Courtesy of the Geography and Map Division, Library of Congress.

No doubt, boxes, and arrows are useful to think about complex organizations in aggregate, “macro” terms. They hide underlying complexity and allow designers to focus their attention at most a few layers deep into the organization. However, typical redesign efforts occur at lower levels in the hierarchy and in smaller portions of the organization. They also require more direct levers to influence individual behaviour and must deal with the gaps between the prescribed formal structure and the actual, realized structure of interactions. As re-organizations of the boxes and arrows amount to different ways to shuffle the hierarchy, such an approach comes into the game with its hands tied when confronting new forms of organization that try to eliminate hierarchy or at least flatten it dramatically.

To give you a sense of how coarse organization charts are, consider the following example: given these four individuals who belong to two departments (R&D – yellow and manufacturing- red), how many organization charts can we create with them? We will use arrows to indicate who gets assigned to work closely with whom. Perhaps you immediately thought of these two:

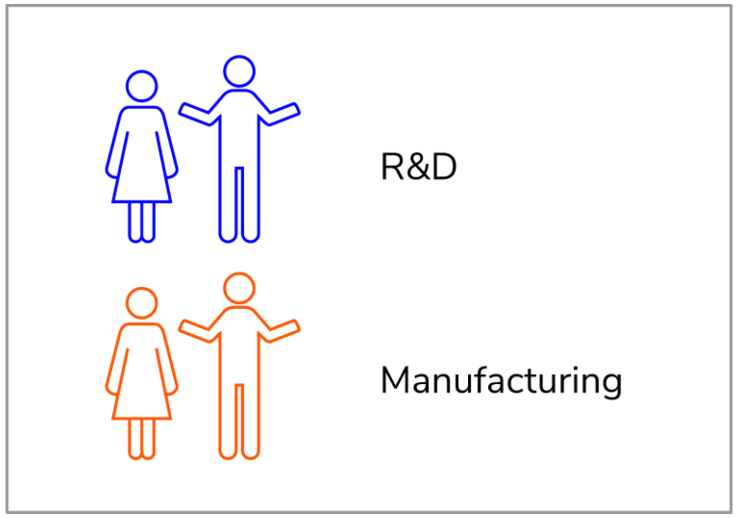

The one on the left is a *functional* grouping, and the one on the right is a *divisional* grouping. Perhaps you also thought of the matrix structure (with both vertical and horizontal sets of arrows in the same picture). But how about any of these?

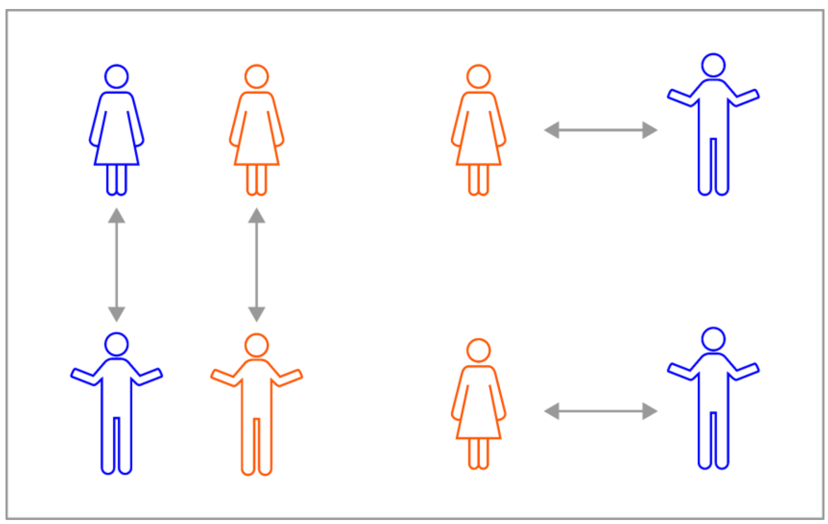

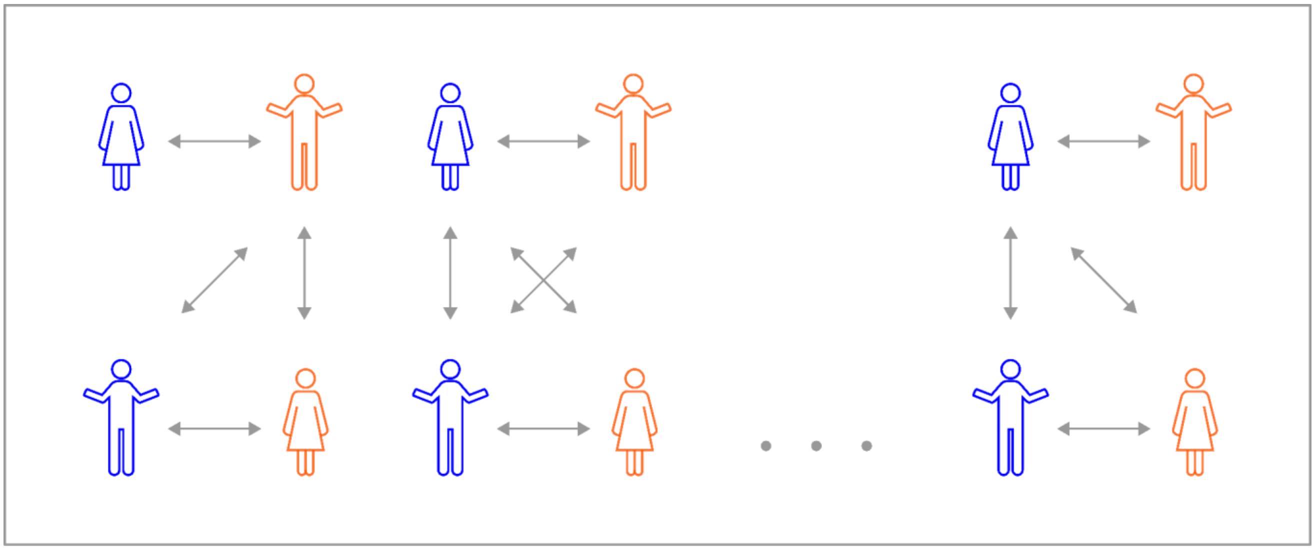

In fact, how many of these alternatives are there? It turns out we know exactly, through the simple application of the binomial theorem: for N unique individuals, the number of distinct interaction structures (i.e. organization structures) is $2^{[(N)(N-1)/2]}$. In our example with N=4, that means 64 different structures. Yes, you read that correctly: the three common organization charts we are used to seeing represent 3 out of 64 possibilities. This brutal simplification depends on what we call the "assumption of uniformity"- i.e., every individual within a function (of the same colour in the picture above) is identical to every other, including in how they interact with other individuals in the organization. Sometimes that is a reasonable assumption; sometimes it is not.

So, is the alternative to consider every possible interaction structure between everyone? The official as well as the unofficial ones? Who listens to whom; who takes orders from whom; who has a crush on whom?

Fortunately, these are not the only options. New thinking about "micro-structures" suggests that we should look at large complex organizations as collections of smaller,

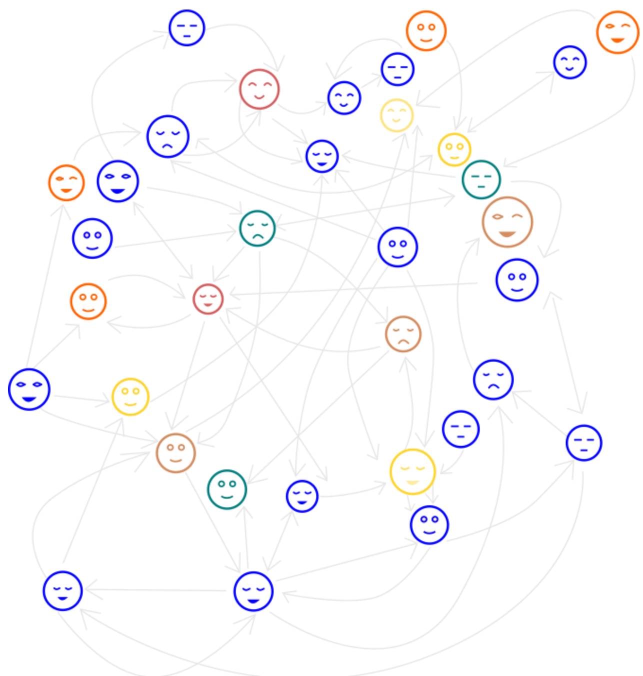

simpler, and recurring patterns of interaction.[4] Start with the simplest possible organization—a pair of people; let's call them Ann and Bob. They can have two types of interactions: ***Interdependence*** and ***Influence***. Interdependence depends on how Ann and Bob are incentivized. If Ann gets paid regardless of what Bob does and vice versa, there is no interdependence between them. But if Ann's outcomes depend on Bob's actions, they are interdependent. Workflows are the most obvious source of interdependence in organizations; if Ann's work depends on Bob's inputs (and Ann's outcomes such as evaluation, bonus, and pay depend on getting her job done), then clearly Ann is dependent on Bob.

Influence implies that Ann can affect Bob's actions (e.g. "She convinced me I should take on this role"). Strong influence can arise from a direct peer-to-peer connection, or indirectly because both Ann and Bob report to a common superior. We will use the same single-headed vs. doubled-headed arrows to distinguish one-sided vs. mutual influence.

We can use this notation, with single and double-headed arrows to distinguish between influence and interdependence, whether symmetric or asymmetric:

[4] Puranam (2018)

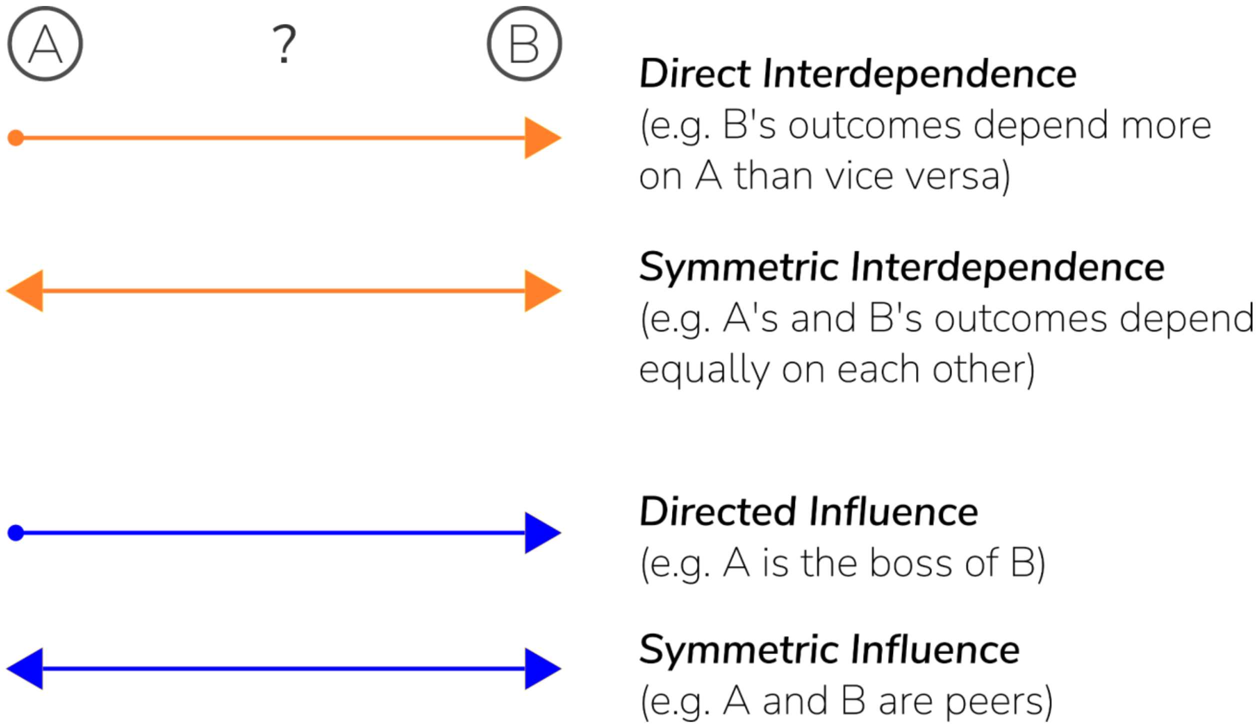

On their own, neither influence nor interdependence is good nor bad. **What matters is the match:** in fact, one of the most basic principles in organization design, also sometimes called ***"the folk theorem"*** (because many scholars independently converged on it with no clear first) is that **influence must match interdependence.** The idea is quite intuitive: if there is interdependence, you need the influence to help manage it. This is at the heart of ideas about "fit" in organization design- the influence structure must fit the interdependence structure, and the latter is dependent on the goals and environment of the organization.

THE "FOLK THEOREM" OF ORGANIZATION DESIGN

The MICRO-Structure Version

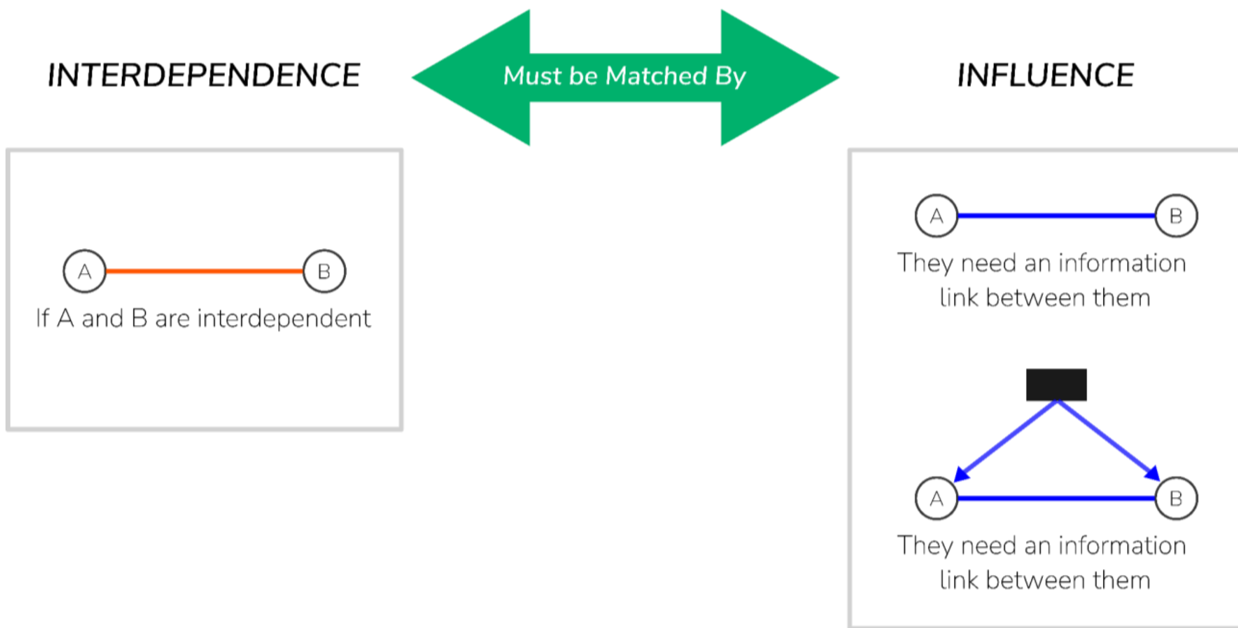

In fact, we can go a step further, to ask **which type of influence should we use to match the interdependence:** symmetric or asymmetric? A useful heuristic is as follows: if the interdependence involves a **"win-win"** situation – both actor A and actor B benefit by aligning their actions- then a symmetric influence tie is sufficient. If however the interdependence is of the **"win-lose"** type (which may still be a net win for the organization), then you might need some form of asymmetric influence- either A influences B more than the other way around (A B in our notation above) or both are subject to a common source of authority (A C B).

Besides changing the patterns of influence to match the pattern of interdependence, one can also do the opposite- **change the interdependence to match the existing influence**

structure - for instance by changing task division and allocation so that A and B no longer depend on each other, or by changing the incentives for A and B, so that their compensation only depends on their individual contribution.

FROM A PAIR TO NETWORKS

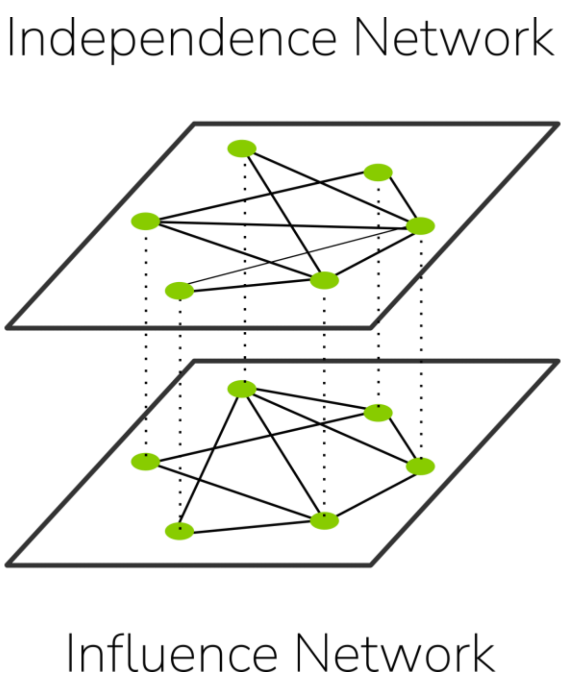

So far, we have focused on a pair – Ann and Bob to understand the folk theorem. But the N-actor picture of the "folk theorem" is not that different, just more complex looking:

The basic design principle does not change; you want to make sure that interdependencies are matched by realized influence network. In fact, every pair of links in the second picture follows the same principle of matching. If we have data on networks of interdependence (e.g. based on business processes or work-flows) as well as networks of influence (e.g. who report to the same boss, or seek advice from each other), then in fact we can use algorithms to do a check for every pair of actors whether the matching of interdependence to influence holds or not.

An important aspect to bear in mind is that the **influence network can be a consequence of both the formal structure (e.g. reporting, grouping) as well as the informal**

structure (e.g. friendship, advise seeking), so a more complete depiction of these various networks might look like this:

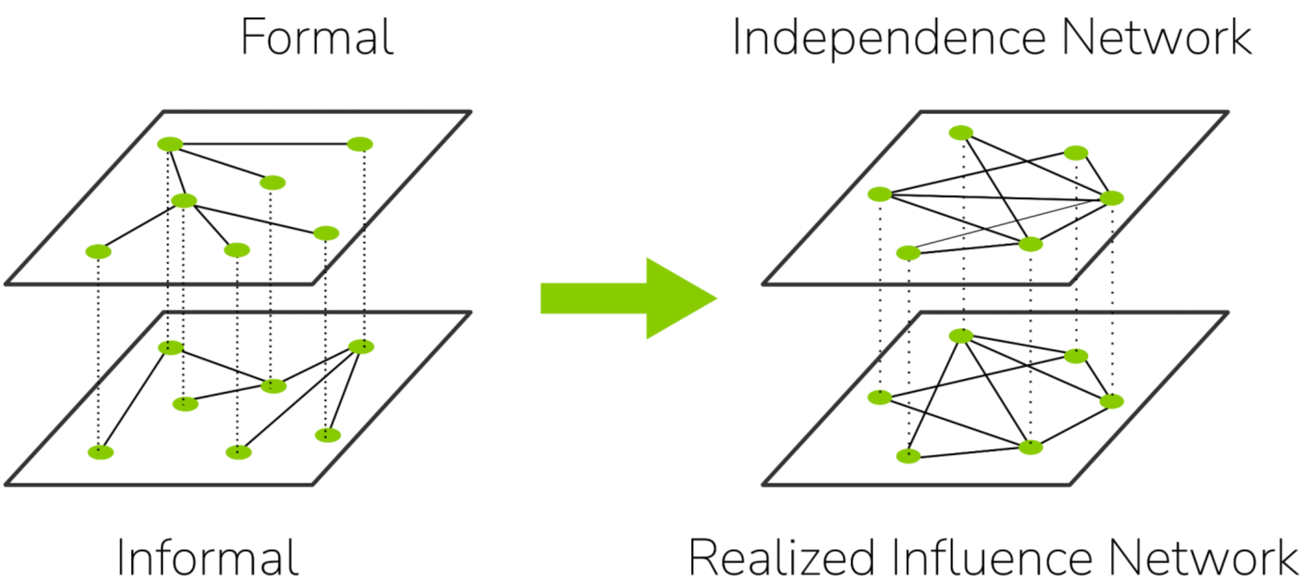

So, here's how we apply the folk theorem or the principle of fit at large scale:

1. we must gather and aggregate data on the formal and informal structure to create the influence network and
2. compare the result to the interdependence network.

As we compare the interdependence and influence networks, we will discover broadly three different kinds of mismatches:

- ***Commission errors*:** an influence tie exists even though there is no interdependence tie
- ***Omission errors*:** no influence ties exist even though there is an interdependence tie

- ***Directional mismatch*:** the direction/symmetry of influence and interdependence ties do not match: A depends on B, but B may influence A

Remember, we do not have to do this tie-by-tie search manually! We have computer programs that can do this in the twinkling of an eye (with the twinkle having to last slightly longer as network sizes grow…. but a twinkle nevertheless!). The detected errors can be corrected at the individual level, and in fact, we might use algorithms to find the most high-impact interventions. Patterns that show omission or commission errors that systematically lie between two departments or geographies may highlight the need for major re-grouping (e.g. divisional to functional structures).

FROM INDIVIDUALS TO GROUPS

Going straight from the dyadic level ("Ann and Bob") to the organizational level makes a network look a lot more complex. But we can simplify that picture thanks to what we call "micro-structures". This exploits a remarkable property of organizational structure: it looks and works in roughly the same way when seen at any level in the hierarchy, or when we replace individuals with entire organizational units. You might say this is because structure is **"*fractal*"**- the patterns look similar at different levels of aggregation. For instance, the basic pattern of a supervisor with multiple subordinates is seen both in the relationships between a CEO and divisional heads, between the HQ of a holding company and multiple affiliate companies, as well as between a foreman and factory workers. Peer-to-peer relationships exist between teammates, divisions, and between companies in a strategic alliance. The clustering of informal friendship ties into clusters within clusters has the same basic structure at any level of aggregation.[5]

These similarities across levels exist because, in organizations, the key **interactions are often between individuals who represent aggregates of other individuals** (over whom they may have authority) This allows us to exploit the fractal property- that at any level in an organization, similar patterns of interaction will have similar dynamics – whether the pattern occurs on the shop-floor or in the boardroom, between individuals who represent other individuals (e.g. department heads) or just themselves (i.e. workers).

So how do we go about re-designing the structure using the ideas we have introduced so far?

---

5 Ravasz and Barabasi, 2003; Puranam, 2018

The **first step** is to **identify the key actors, and the interactions among these, that you think are at the heart of the design problem you are trying to solve.** For instance, if you are trying to correct a tendency for certain decisions to be slowed down because they all pile up at one place in the organization(a.k.a. the ***bottleneck problem****)*, identify the key roles/units involved, and map out the interactions between them. Who needs decisions to be approved? Who has the authority to approve which parts and kinds of decisions? Where do the delays typically occur? Or perhaps, you see frequent ***breakdowns*** in collaboration or missed opportunities across departments. Which units or roles should be collaborating but are not? Again, you should identify the key roles/units involved and draw how they interact. Or your issue may be the classic problem of the wrong persons in the role, who doesn't have the right skills and so creates bottlenecks and breakdowns (a.k.a. The **"*bumbler" problem***). Who is affected by the poor decisions of the bumbler? How are they connected to each other? Yet again, you need to identify the key actors and the pattern of interactions among them.

The second step is to distinguish clearly between two different types of interactions: interdependence or influence. You can use workflow data and process mapping to understand the interdependence structure. Organizational charts, surveys, and seating plans can give us a picture of the influence structure. Increasingly we can also use digital exhaust from emails, video calls, and collaborative chat tools (like Slack/Teams) for the same purposes.

The **third step is to apply the principle of fit (folk theorem):** find ways to improve the match between interdependence and influence. Remember this can be done by changing influence structures, interdependence structures, or both.

Being able to reduce the complexities of organization design problems to a picture with the smallest number of actors and their interactions is a skill. But the good news is it is a skill that you can build with practice. Here is a case to get you started:

Use Case#1: Re-designing to eliminate bottlenecks

Slow decision-making because of bottlenecks (i.e. places in the organization that function like black holes when it comes to decisions) is a common signal for initiating an organizational re-design. Slow decisions mean cost overruns, missed opportunities, and ultimately competitive disadvantage. So how would we solve this problem solving Org2.0 techniques?

STEP 1: Perception

The first step, as always in analytics-based design is Perception. We gather data on the existence and extent of the problem we are trying to solve. How do we define a bottleneck? Who are the key individuals involved in the decision process? For the key individuals, can we identify the interdependence and influence relationships among them? How much must a decision be delayed for us to consider a bottleneck to exist? What is the (opportunity) cost of the bottleneck? Is there really a bottleneck or are we just hearing subordinates complain about superiors who rejected their proposals (perhaps on good grounds)? If a proposal was not delayed but rejected, what is the evidence that this was in fact a bad decision?

One of the first things we are likely to learn from the perception phase is how widespread the problem is: is it ***local or global***? If it just involves one individual, then we may be looking at a sorting problem (i.e. finding better talent for this role). We address sorting problems in the next section. It could also just be a simple capacity problem, in the sense that the individual has no bandwidth to deliver timely decisions (perhaps they have too lager a span of control or do not have clarity of responsibility and delegation), in which case we can just make the needed tweak to the structure for this localized problem.

But if the same pattern of bottlenecks, with the same or similar set of roles (e.g. managers from sales, manufacturing, and somebody in the top leadership team) always show up as the cast of characters, then we know we are possibly looking at something more global than local, that might require a significant re-design. We proceed here on the assumption that this is the case- that we are seeing a pattern of delayed decision-making involving not necessarily the same specific individuals, but perhaps the same sets of roles in different parts of the organization.

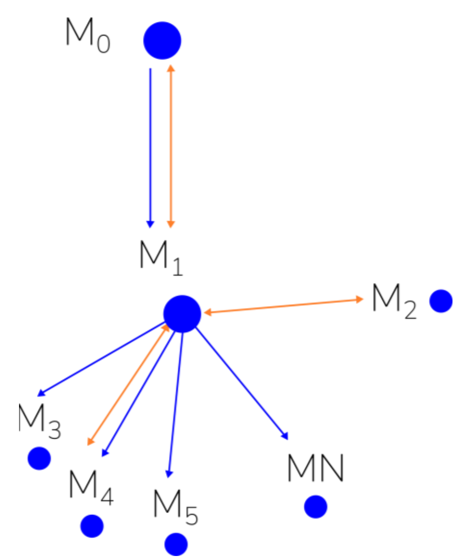

Figure 1 illustrates the figures we might draw to understand the bottleneck problem. Decisions that M_4 needs from M_1 seem to be the ones that get held up forever. It could be that M_1 is overburdened; has insufficient delegation from M_0 or is struggling to manage the interdependence with M_2. There are many possible solutions, which could involve changing the interdependence structure (e.g. by changing incentives or task

allocation) or the influence structure (to create a new reporting or coordinating relationship). Which one will work will require knowledge of the context. These "micro-structure" diagrams help with diagnosis and a thorough consideration of alternatives.

STEP 2: Prediction

If the problem is indeed widespread and recurring, this also makes it possible to gather data at a large scale on the decisions that have been bottlenecked vs. those that have not. This is where Prediction based analytics comes in useful. Can we forecast which decisions are most likely to end up in the "black hole"? Are there certain attributes of those decisions (e.g. the ones originating in a particular department, involving a particular set of issues, or a level of financial implications) that make it more or less likely to be bottlenecked? With prediction-based analytics, we can take one of two approaches: we might build the "early warning system", that tells us which decisions can be expected to be stuck, so that we can spend our limited resources and energy just chasing those to make sure they don't. In such an approach, we would be agnostic to the "why" question-the reasons that make certain decisions more or less likely to be stuck. Instead, we are pragmatic and focus on preparing to deal with the symptoms before they arise, based on prediction models that can make such forecasts. The use of "black box" (hard to interpret) algorithms may be perfectly OK for such an application.

But an alternate approach is to go for "root cause detection"- to understand why certain decisions seem more at risk of getting stuck (and then to do something about that). This will usually limit the predictive analytics models we can use to those that are relatively easier to interpret. But the good news is that such models are also precisely the ones that do not need massive amounts of data. In practice, this is the approach we think is most likely to form the predictive analytics strategy for the bottleneck problem. Using such an approach we may see that decisions that involve a particular kind of complexity (requires consideration of the impact on 2 or more departments) that arise in the 3rd financial quarter of the year are the ones that are at the highest risk of getting delayed (to take an instance). With this insight, we might now want to think of how to solve the problem by changing how the organization processes these decisions. But remember, predictive analytics only generates indicators, not interventions! The results at this stage are correlational so we still need to prototype before rolling out the changes.

STEP 3: Prototyping

Let's suppose you believe that adding a coordination meeting between managers M1 and M2 in Figure 1 would help speed up the decision-making process of the cases that have an impact on both groups. How do we know this will solve the problem? Prototyping is all about knowing whether your proposed intervention will cause an improvement. The way to assess that is to think in terms of counterfactuals –what might happen if all else stayed the same except your intervention was implemented. As we noted in Chapter 3, there are a range of techniques one can use to prototype, ranging from scenario-based discussions (e.g. getting a group of experienced managers to think through the last few decisions that were bottlenecked, and doing the thought experiment- would your proposed change have prevented the problem?), to computer simulation (e.g. modeling the flow of decisions and time taken for each on average to understand how frequent and effective the coordination meetings would have to be to ease the bottlenecks), and culminating perhaps in a field experiment (e.g. try the new system in some parts of the organization, create a control group where it is not implemented and compare the average speed of decision making in both).

Which one you want to use will depend very much on the urgency of the problem and the resources at your disposal. But it will also depend on the ***value of information***- it is only when both the potential upside from the intervention working is huge, as well as the potential downside of it misfiring is huge that you need elaborate prototyping. If at least one of these conditions does not hold, you don't need to prototype- you would either not bother implementing the intervention or implement it anyway without waiting. We also reiterate that even something as simple as scenario analysis can already help to stop bad interventions in their tracks. For instance, experienced managers might be able to point out to you that M_1 and M_2 will be very unlikely to agree in their coordination meetings, because their interests are not aligned, and their first common boss is several layers away. In this case, we need to go back to the drawing board and think of other interventions.

Reflection question: Compare the approach you undertook the last time you had to re-design something in your organization, with the Org2.0 approach we prescribed above. What are the barriers to using our approach instead?

CHAPTER 5

The second lever of organization design- Sorting

Overview

- Sorting refers to choices that determine who works in an organization. Two organizations with the same structure but different people will perform quite differently.
- Improving hiring and retention through data has become one of the most successful applications of organizational analytics. However, sample selection bias remains a key challenge- we often only have data on those we hired, and this can mislead us about whom we should hire.
- Prototyping through A/B tests or even some degree of "hiring at random" can be a powerful antidote to biased data.

Expanding our thinking about structure is useful, but we should be aware that structure is only one of the approaches to solving organization design problems. S*orting* is another.[6] When using structure as a design lever, we take the composition of individuals in the

[6] Some of the seminal work on non-structural aspects of organization design has been reviewed under the broad rubric of "culture" in Schein 1985; O'Reilly and Chatman, 1996; Chatman and O'Reilly, 2016.

organization as a given, but sorting does not. Sorting can occur in both directions- into an organization as well as out of it. Further, it can occur within an organization in particular projects. groups and teams.

As organizations are systems with boundaries, the entry (and exit) decision is a critical one. Individual-level variations in goals, skills, and beliefs can be exploited to create better matches between individual attributes and organizational requirements in terms of skills and goals. A particular form of sorting is self-selection- in which individuals sort themselves into an organization, and even within the organization, into particular tasks. This has the advantage of generating intrinsic motivation through the freedom to choose as well as the match between what is chosen and what the individual prefers.

Consider the following: Team A was formed by a group of individuals who self-selected themselves into it because they were deeply motivated by the goals of the team and the prospect of working with like-minded individuals. Team B was composed of individuals who were assigned to it because they were available. In which team do you think motivation and commitment levels will be higher? In which team will colleagues be more often willing to make adjustments and resolve issues peer-to-peer? Which team do you think is more at risk of getting trapped by groupthink? Which team is more likely to be cliquey and have lower diversity? As you can see, even if Team A and B had the same objectives and the same formal structure, the way they work would be quite different.

As we often like pointing out ***employees are not randomly assigned to different organizational structures.*** They are sorted in by themselves and of course by HR professionals. No wonder, structures that appear to work with one set of individuals often do poorly with others, and vice versa. In fact, this is one of the reasons why imitating industry best practice is dubious – a structure that works for Company A given who works there is unlikely to work for Company B if their employees are different.

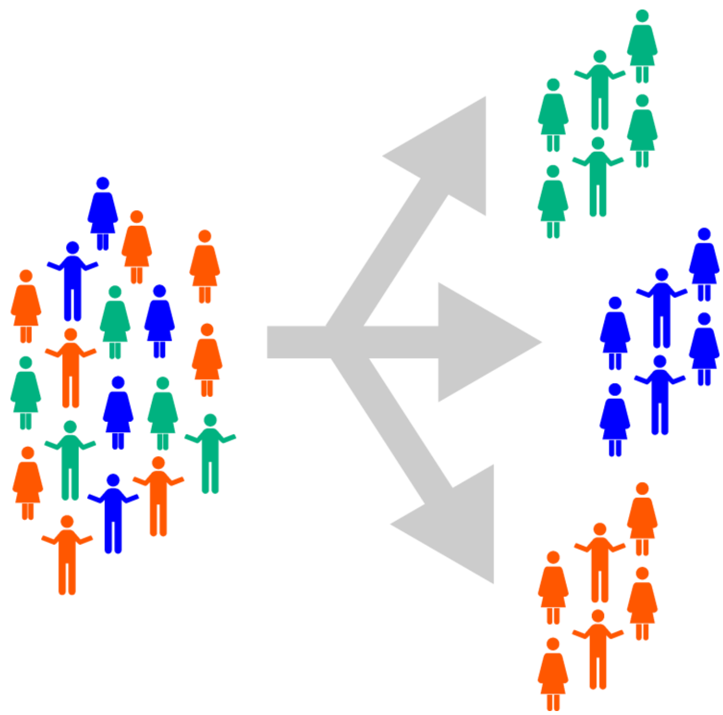

The application of data analytics to improve hiring, retention, and staffing decisions first emerged as a widely popular instance of Org2.0 thinking. Traditionally, HR professionals have relied on their past experience and judgement to make decisions on whom to hire and whom to try to retain based on considerations of fit. But increasingly, we can use data and algorithms to possibly do better. **To be clear all decisions are always based on data and algorithms, though sometimes, the algorithms sit between our ears, and the data is called "experience".**

The question simply is whether we can do demonstrably better by using electronic databases and computer algorithms.

Use Case#2: Improving diversity in the talent pipeline

Sorting is pretty much the "baseball" of organizational analytics: it's the area in which analytics are the most developed, mainly because it's easy to identify the unit of analysis (people) and we don't need to consider as many complex interdependencies in understanding what makes people give their best as we would need in understanding what makes a whole team or unit work best.

Many applications have been around for some time, but in recent times one issue has been ruling them all: diversity. Organizations have come under great pressure to employ a workforce more representative of the broader society. That means not only hiring a diverse workforce but also making sure that diversity remains intact from the bottom to the top of the organization by providing equal opportunities for promotion and development to all populations. The lack of diversity can have different sources and managing opportunities within an organization that already exhibits low diversity can be difficult.

Few organizations apply sophisticated methods to this problem. The typical case goes along these lines: top-level managers recognize that the organization has a diversity problem when an embarrassing statistic or a prominent case of discrimination makes the press, and they identify areas of the organization that seem most guilty of excluding minorities by simply looking at which unit present the least demographic diversity. They might scramble to hire minorities in those areas (or promote minorities if the lack of diversity is only prominent at the top of the organization) without much consideration for whether diversity issues might show up again in the future and whether this detracts from the organization's ability to function well (hiring minority managers on an emergency basis is a lot less effective than preparing employees for managerial roles far in advance). Can we do better than this? Applying Org2.0 thinking should help organizations scan for a lack of inclusion in their workforce over time, identify specific minority employees at risk of leaving the organization, understand why they might be leaving and identify policies that can help hire, retain, and promote them.

Let's look at each step in the organizational design analytics cycle.

STEP 1: Perception

A good analytics cycle starts with good data. People provide information about themselves when they apply for a job; they go through an application process that leaves (or should leave) digital traces at its different steps; and legal matters make it necessary to gather data about when employees start their job, what that job is, and how well the employee has been performing in that job. Keeping track of that data should be common sense, but it is definitely worth emphasizing: many companies have trouble even knowing how many employees work for them. There are always immediate concerns that seem more urgent than creating a good data infrastructure, but in the long run, it definitely pays off to gather workforce data systematically. Making the effort to link different databases together is worth it, too: by merging data about applications, job interviews, work performance and promotions, you will already start to formulate hypotheses about where your company's lack of diversity comes from. It might stem from not attracting a diverse enough applicant pool, from weeding out minority applicants through a biased interview process, or from reducing progression opportunities for the employees you do hire (because of a negative environment or a biased performance appraisal and promotion process). Recognizing which stage is the most problematic (and how this may differ in different units of your organization) is already valuable.

STEP 2: Prediction

A predictive approach to the talent pipeline is helpful both in order to identify specific employees at risk of leaving and understanding the factors which put them at risk of doing so. The easiest application would be an early warning system identifying employees at risk of leaving the organization in the near future. The approach is simple: collect all the data you may think of about your employees, generate as many variables as possible based on that data (about employee demographics, performance, working environment), apply different machine learning algorithms to the data you collected in the past to see which algorithm makes better predictions, and finally make predictions about the future using the algorithm which seems to perform best. The whole process shouldn't take more than a few days for a good analyst with clean data and will leave you with precise predictions about which minority employees seem the most at risk to leave the organization. You might want to at least talk to these employees to understand whether and why they might be thinking of leaving, and whether there is something you can do about it. In the same spirit, you could use your employee data to predict which managers are the most likely to promote minority employees in the future.

The approach above is completely agnostic to why employees might be leaving: all you care about is that they are likely to leave and that you know about it early enough to inquire about them. But knowing why minority employees tend to leave can be very valuable as well and can generate ideas about how you may provide them with a better working environment. For instance, you might find that minority employees are especially likely to leave the organization when they rate their managers poorly. Perhaps providing some training to your managers to sensitize them about diversity can be useful? You might find that minority employees who went on voluntary team-building trips are much less likely to quit in the next year. Maybe making these trips more widely possible could help? Whatever policies you come up with are only a starting point, because machine learning cannot tell you whether any variable actually *causes* employees to exit (maybe those employees who went on team-building trips were simply more committed to the company, to begin with!). But you can take your policy ideas to the next step: prototyping.

STEP 3: Prototyping

Prototyping is all about generating counterfactuals. By counterfactual, we mean "what would have happened if X had not been true". "X" can be anything, for instance, "going on a team-building trip" in our example above. Maybe most employees who went on team-building trips didn't quit in the next year; but would they have quit if they hadn't gone? Chapter 3 described several sophisticated ways of generating counterfactuals. But simply thinking about counterfactuals is already a great start. You may get a few managers to sit together and re-examine past exit interviews, for example. For each minority employee who left in the last year(s), does it look like organizing more regular feedback meetings with their manager could have helped keep them committed to the company? A similar approach can be applied to promotion processes: for each recent promotion case that did not yield a minority candidate, could we have found equally qualified minority candidates by involving more managers in the search process? Scenario-based thinking is a very powerful and inexpensive tool.

Combining insights from predictive analytics and scenario-based thinking, you should be able to come up with a few ideas for policies that might help alleviate diversity issues. But how do you know they will work? Do you need to roll them out in the whole organization to figure that out? Often, you don't! For instance, you may think that part of the issue comes from minority candidates not applying for managerial positions. What if, instead of waiting for applications, you sent targeted emails to candidates whose qualifications look like they fit a position? You may start by doing this on a random basis for half of the positions you seek to fill in the next year, and then roll out the policy for all promotion cases if results convince you that this is good practice.

We close this chapter with a warning: it's also important to be aware that hiring data has some unique issues that can make its unthinking use quite dangerous. Here is the core problem: typically, we observe the performance of people we hired, not those who we did not. Therefore, using data on current employees to make guesses about whom to hire is prone to what is known as the ***sample selection problem.*** Prototyping your interventions de-risks this. Surprisingly, so does hiring a small portion of your workforce at random! This article below explains why.

Article on biases in data and how randomness can help beat them

EMBEDDED LINK

Reflection question: Using data on employees to improve organizational performance is not necessarily the same as improving employee welfare. When are these objectives aligned? What ethical safeguards (besides of course *always* complying with regulation) should you put in place when these are not aligned?

CHAPTER 6

The third lever of organization design- Sensemaking

Overview

- How people make sense of their organizational context – as reflected in things like culture, identity, vision, and leadership- affects how the organization performs.
- Two organizations with identical structures and similar workforces may nonetheless perform differently because of differences in these sensemaking processes.
- Analytics can help in assessing the similarity of these processes across organizations (e.g. assessing cultural fit in mergers) as well as the effectiveness of interventions aimed to change these (e.g. forecasting if organizational change projects will take off or sputter out).

Sensemaking refers to the process through which shared meaning is constructed within an organization as its member create a mutual understanding of experiences and situations.[7] Framing an issue, describing a vision, and leadership are processes of ***one-to-many*** sensemaking; a consensus formation is a ***many-to-many*** form of sensemaking. In either case, neither the set of individuals nor their patterns of interaction need to change,

[7] Weick, 1995.

but there can be a significant change in the individual's beliefs and goals over time, leading to changes in behaviour. Org2.0 thinking relies on modeling such processes rigorously using computational and network tools.

Culture analytics: the Org2.0 approach

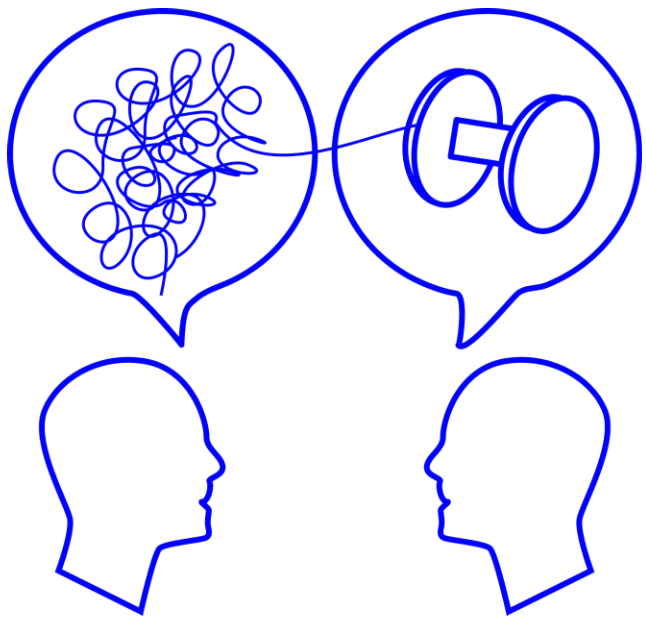

How can we measure an organization's culture? Without question, culture is important- it refers to the set of widely and deeply held mental constructs, such as values, beliefs, and norms. It shapes what employees do when managers are not looking. But how can we measure, let alone shape it? One of the major advances in recent years has been the application of machine learning to text- in the form of survey responses, chat, emails, or employee descriptions of their company (such as are found in job posting/employer review websites). One can not only describe an organization's culture using such tools but also compare cultures across organizations. That can be very useful in situations involving collaboration between organizations, such as alliances, mergers, or even between divisions within a company. One can gauge the sentiment and sense of engagement of employees, and isolate trending issues and themes.

An article outlining how text analytics helps to understand organizational culture

EMBEDDED LINK

Use Case#3: Assessing cultural compatibility using Natural Language Processing

It's a common argument among managers that culture is too complex to apply analytics to. Sensemaking is about what people think; it's about which values they hold and how they go about communicating those values. Surely the only valid way to understand this is to ask them about it? Not really. Taking the time to talk with your employees and understand what's on their minds is a great idea in theory, but you won't be able to ask each and every one of them about what they think. And they might not tell you to your face what they might be willing to say anonymously. What if you could get insights into what all of your employees are thinking? Recent developments in text analysis, combined with new sources of information about employees' perceptions of their companies, allow you to generate those insights at scale.

STEP 1: Perception

One of the major advances in recent years has been the application of machine learning to text. The idea behind one of the most widely used approaches known as 'topic modelling" is simple. When employees write about their company, they have certain topics in mind which can be described using specific words (for instance, we typically do not use the same words to describe work/life balance and compensation). Topic modelling is about uncovering the topics that are important to writers by looking at the words they use when describing them. You could collect this data from a variety of sources, from internal surveys to emails/chat archives and even external job posting/employer review websites. To tell you the truth, someone probably already has done that work for you:

some prominent consultancies now routinely collect data (for free!) from websites such as LinkedIn and Glassdoor to better understand their clients. Simply looking at Glassdoor reviews written by your employees tells you a lot about what your employees care about (for instance, is compensation one of the topics they mention the most? Or do they seem to care more about doing challenging work?), how they think the company is doing on those aspects (do they write about compensation using positive or negative terms?), and how these values are distributed across the company (does everyone seem to care about the same few things or is there wide variation? Are there systematic differences between different units in the organization?). You will also be able to track the evolution of your company's culture by collecting this kind of data over time.

STEP 2: Prediction

Simply being able to observe what your employees care about is already very valuable; it allows you to take the pulse of your organization and check whether its culture is evolving in a direction that aligns with the culture its leaders are hoping to instill. But you can even go further and use these insights to make predictions about the future of your organization. Many aspects of an organization's future are affected by its internal culture: whether it will successfully integrate new employees, how successful it may be in collaborating with other organizations, and even how well it may execute a strategy. We will focus on one example which we think is quite promising: the management of mergers between companies. Mergers are a very common occurrence in today's business world, and it's widely acknowledged that they also very commonly fail. Culture is a major reason for this: mergers look like great ideas on paper when they offer opportunities for synergies, but synergies typically require collaboration among the two firms' employees. Cultural differences make collaboration harder: people have a much easier time working with people who share similar values and communicate in similar ways.

Can we predict ***cultural clashes*** in mergers? We can, using text analysis[8]. Glassdoor, for example, provides reviews about thousands of companies. The odds are pretty high that you can find reviews from the company you are thinking of acquiring! You can easily compare the topics most prominent in two companies to better understand their respective culture, which should help you predict potential conflict areas (for instance, employees who care mostly about challenging work may experience some frustrations working with colleagues who mostly care about work/life balance). Making those predictions will require some judgement on your part (most tools won't reliably tell you that your merger has an X% probability of failing because of cultural issues). But text analysis provides a very reliable starting point for understanding cultural differences: when you spot significant differences through text analysis, you might invest then in further information (surveys, interviews) to understand whether they reflect true cultural differences between companies, and use common sense to predict their consequences. In the extreme case, this might lead you to renounce a merger deal. In most cases, it will at least provide you with valuable insights about which cultural differences to watch out for during the post-merger integration phase.

STEP3: Prototyping

Websites that store employee reviews typically do so for a large number of companies. So why would you only analyse the culture of your one preferred acquisition target? You

[8] Marchetti (2020).

can browse through as many possible targets as you would like to see how their respective cultures differ from yours. At the very least, this will provide you with benchmarks to understand just how different your preferred target's culture is relative to other potential targets. But it might also change your opinion about which target should be the preferred one in the first place.

Change management in the Org2.0 way

Organizations change things all the time; at least they try to. Entering a new market, changing an organizational structure, implementing new enterprise software...all are examples of major organizational change which can be decided by top managers but require buy-in across an entire organization to be truly successful. Many managers consider change management as the most difficult aspect of their job. We all know there is comfort in preserving the status quo, and most change efforts encounter stark resistance from employees. So, how can you get buy-in? There is a lot of received wisdom on the topic, typically around the importance of early buy-in by informal opinion leaders: the idea is that once you convince your most central employees about the value of change, they will go around and convince others until everyone believes in it. The reality is a lot more complex, and it's unlikely you will be able to predict it correctly without a systematic approach.

Use Case#4: Forecasting change adoption using computer simulation

STEP 1: Perception

Broadly, change management depends on two things: what your employees believe about the change you are proposing, and how they can influence each other to change those beliefs. The first step is to generate a reasonably accurate picture of both.

Interviews and surveys can be useful for understanding your employees' attitudes toward change. Interviews are especially useful because people might not feel comfortable sharing their honest opinion in surveys. It can be a good idea to ask your most important managers what they think about a change initiative, and also ask them what they hear others around them think about it. As for understanding how they can influence each other, what you will need is to form an accurate picture of your organization's social network: who talks to whom? There are different ways to get data on this: depending on what your employees are comfortable with, you may either ask them directly about who they talk to (surveys) or use unobtrusive measures such as email traffic to better understand who talks to whom. Simply looking at a combined picture of the social network and the attitudes of employees towards change should already give you a few hunches about your likelihood of success: if your most influential employees are against the change, you are probably not standing on very solid ground. But we can do even better if we could figure out what the most effective interventions might be to increase the chances of successful diffusion of change.

STEP 2: Prediction

The diffusion of opinions in an organization is a complex process. Because it depends on collective dynamics among many different people who influence each other, it is quite hard to predict what the outcome will be: will the change progressively convince more and more members, or will it end up being rejected even by those who started out in favour of it? To make things even worse, organizational change is a relatively rare event: organizational transformations don't happen every week, and even in tumultuous organizations today's change is rarely perfectly comparable to yesterday's. This means that organizational analysts have little prior data to learn from when predicting the success of the change. Does that mean we have to fall back only on intuition? No! Just because we can't *perfectly* predict the outcome of a complex process doesn't mean we shouldn't try to predict it *as well as we can*. In the absence of past data, what we rely on is theory: we try to get an accurate picture of the different pieces of the puzzle (in this case, people's opinions and the social network that links them) and try our best to think of how these pieces interact together to produce the outcome we're interested in predicting.

Luckily, we don't need to rely exclusively on our brains to understand these complex interactions. We can leverage the power of computers to simulate the process of organizational change through what is called **agent-based models**. In these models, we represent an entire organization (for instance, we can represent its different members, what they believe, and who talks to whom) and set assumptions about how the organization evolves over time. For instance, we might assume that every week, any two people who know each other can talk to each other about the change initiative and one might convince the other. By simulating the model over time, we can get predictions about the likelihood of successful change implementation. These predictions answer the question "Given what we think we know about our organization, and given how we think people influence each other over time, what is the likelihood that most of our employees will buy into the change?". The approach is actually quite similar to how the spread of pandemics is modelled. If we know something about the structure of connections between people and their infectiousness and susceptibility to infection, then we can make a pretty good guess about how many victims we will have on our hands. The spread of ideas is similar though not identical: it may require multiple exposures before an idea is accepted and unlike resistance to disease, opposition to an idea can also be infectious.

STEP 3: Prototyping

The beauty of computer simulations is that they also allow us to experiment with different scenarios *in silico* (see Chapter 3). Since the analyst fully controls the simulation parameters, it is easy to create counterfactual scenarios where only one variable is changed: "What is the likelihood that the change will diffuse throughout the organization if we don't do anything? What about if we convinced person X about the change? What about if we convinced person X and named her 'Change Initiative Officer'?" What's even more beautiful? You are running these experiments on the computer, practically for free. This is similar to checking the sensitivity of your financial projections to different scenarios using a spreadsheet. With good models, we can do the same for organizational processes.

An article outlining how computer simulation helps test our intuition about when framing and vision may help collaboration

EMBEDDED LINK

Reflection question: How do you currently make decisions about assessing culture or managing change? Are you intuitively following the Org2.0 approach but without data, or do you do something altogether different?

CHAPTER 7

An Eye on the Future

This book is <u>not</u> about the latest organization designs – it is about a new approach to design*ing* organizations. By this point, we hope you have developed an appreciation for organization design thinking the Org2.0 way- **please don't adopt a new organization design without thinking carefully about whether it can work for <u>your</u> context**. Data and algorithms can help with this—that's the basic message of this book.

In this book, we already gave several use cases for how to apply Org2.0 techniques. But these are hardly the complete list. For instance, as the world of organizations adjusts to the post-pandemic era, one thing is clear- remote collaboration will feature prominently. Here is a sample of things that Org2.0 techniques may help with as companies prepare to tackle this change:

- Identify how to modularize work as well as prioritize what needs to be done at the office - to create optimal shifts for remote collaboration
- Identify who emerged as "super-connectors" to coordinate across teams virtually
- Diversity and inclusion- is it getting worse or better with remote collaboration?
- Text analysis to identify cultural fragmentation; and identify "cultural champions" who strengthen and reinforce the organization's culture.

How we design organizations in the age of algorithms is itself likely to change, as developments in AI algorithms and innovations in organizing continue to unfold at breakneck pace. In this edition of the book, our focus has been on algorithms as ***design tools for organizations.*** Statistical analysis of data (including inductive inference from machine learning, traditional deductive inference from experiments, and network analysis) relies on algorithms for processing (possibly large quantities of) data. These applications of algorithms are quite concrete—they involve fully specified instructions in the form of computer programs. We also discussed the use of computational algorithms as representations or models of organizations. One can for instance simulate how the division of labour would look under self-selection instead of authority, to understand its relative advantages and disadvantages. We are now developing a generative AI based organization design co-pilot- think ChatGPT but with specialist expertise in organization design. Look out for it!

Another approach that has been gathering steam is to use algorithms as ***embodied solutions*** to organization design problems. For instance, Uber's ride-sharing service uses an algorithm to do task allocation (i.e. rides to drivers), GitLab's "continuous integration" software coordinates the coding efforts of remotely distributed programmers, and ODesk's algorithms compute compensation for tasks based on the nature of those tasks and worker skills. Previously, each of these organization design solutions would have been instantiated in the form of managers making and implementing decisions.

It is certainly possible that in the near future, we will see organizations in which at least some of the ***workers are algorithms***. This goes beyond the use by human agents of algorithms to make decisions—rather we are talking about intelligent agents with some degree of decision rights. How such agents will work in conjunction with human agents, and what this means for the design of organizations are the key questions. For instance, how the division of labor may occur in mixed teams of humans and intelligent algorithms, how they co-specialize, and how their joint productivity can be enhanced are all topics within this theme.

That may well bring us to Org3.0!

Video and article on what the future of organization design may look like

EMBEDDED LINK

References to academic literature cited in text

- **Burton, R. M., & Obel, B. (1984).** *Designing Efficient Organizations: Modeling and Experimentation*.
- **Centola, D., & Baronchelli, A. (2015).** The spontaneous emergence of conventions: An experimental study of cultural evolution. *Proceedings of the National Academy of Sciences of the United States of America*, *112*(7), 1989-1994. doi:10.1073/pnas.1418838112
- **Chatman, J. A., & O'Reilly, C. A. (2016).** Paradigm lost: Reinvigorating the study of organizational culture. *Research in Organizational Behavior*, *36*, 199-224. doi:10.1016/j.riob.2016.11.004
- **Clement, J., Gaba, V. and Puranam P. (2020)** *Organizational change as a complex diffusion process*, Working Paper.
- **Marchetti, A (2020)** *Organizational Culture and Performance*, Ph.D. Dissertation, INSEAD. Mintzberg, H. (1979). *The Structuring of Organizations: A Synthesis of the Research*: Prentice-Hall.
- **Moreland, R. L., Argote, L., & Krishnan, R. (1996).** Socially shared cognition at work: Transactive memory and group performance. In *What's Social about Social Cognition? Research on Socially Shared Cognition in Small Groups* (pp. 57-84): Sage.
- **Puranam, P., Alexy, O., & Reitzig, M. (2014).** What's "new" about new forms of organizing? *Academy of Management Review*, *39*(2), 162-180. doi:10.5465/amr.2011.0436
- **Puranam, P. (2018).** *The microstructure of organizations*: Oxford University Press.

- **Reagans, R., Miron-Spektor, E., & Argote, L. (2016).** Knowledge utilization, coordination, and team performance. *Organization Science, 27*(5), 1108-1124. doi:10.1287/orsc.2016.1078
- **Schein, E. H. (1985).** *Organizational Culture and Leadership*: Wiley.
- **Weick, K. E. (1995).** *Sensemaking in Organizations*: Foundation for Organizational Science-Sage Publication.