Stanford | ENGINEERING
Management Science & Engineering

"Rolling Up the Leaf Node" To New Levels of Analysis: How Algorithmic Decision-Making Changes Roles, Hierarchies, and Org Charts

Melissa Valentine
mav@stanford.edu

Rebecca Hinds
rhinds@stanford.edu

May 6, 2021

Abstract

This paper draws on a 10-month ethnography of a retail technology company to contrast two conflicting ways the organization structured the work of an internal professional group. The original structuring approach involved using *bureaucratic decision structuring* to coordinate the collective decision-making of that group. This structuring set up the profession's organizational chart as a "decision tree"—but one that had an unexamined impact on decisions and outcomes and was not expected to change very much or very often. In contrast, the second structuring approach involved using *algorithmic decision structuring* to coordinate collective decision-making. This structuring approach involved explicitly measuring the impact of decisions on outcomes—including measuring the impact of the overall decision structuring itself. This approach also involved the expectation that the structuring would change frequently and responsively. To use the new algorithmic tools, the organization had to reconfigure not just single workflows, but the larger set of role relationships and hierarchies that were structured by the entire organizational chart. Because organizational charts and bureaucratic decision structuring have been a prevalent mode of production for many decades in many industries, we propose that this difference between approaches to decision structuring offers an undertheorized reason why algorithms bring about large-scale infrastructural change.

INTRODUCTION

Responding to the so-called ‘Big Data revolution’, many scholars have examined how algorithms are reconfiguring professional expertise (e.g., MacKenzie, 2014; Christin, 2018; Sachs, 2019; Brayne and Christin, 2020). However, to date, there are few theoretical accounts considering how data and algorithms are changing the *structuring* of these professions, for example, their role structures, professional hierarchies, or organizational networks (Bailey and Barley, 2020). This gap is surprising because the main theories of technological change and its consequences hold that technologies are occasions for structuring (Barley, 1986) and that technological revolutions are revolutions precisely because, by changing the means of production, they significantly alter how people interact within and between established professions (Barley, 1990; Barley, 2015; Barley, 2020).

To date, the extant literature on algorithms and professional expertise has focused more on how data and quantification are changing the subjective nature of professional expertise. These studies note the juxtaposition of the abstract and almost “sacred” nature of professional expertise against the perceived accountability and objectivity of algorithms (Christin, 2017: 13). Professionals now heed the “rules” of algorithms, changing the perceived value and source of their professional expertise: journalists are incorporating audience preferences in their reporting (Christin, 2018), radiologists are leveraging recommendation for medical diagnoses (Lebovitz, Lifshitz-Assaf, and Levina, 2019), and high-frequency traders are relying on algorithms to make financial trades that balance profitability with low risk (MacKenzie, 2018).

Our paper examines instead how data is changing the *structuring* of professionals’ collective expert work practices, including their role relationships and professional hierarchies. This question builds on research that has examined the implications of professionals being employed in bureaucratic organizations (Huising, 2015; Anteby, Chan, and DiBenigno, 2016), wherein their jobs are nested in roles or "systems of prescribed decision premises" (Simon, 1991). Those professional roles, in turn, are nested in managerial hierarchies and subject to financial and bureaucratic controls that distribute responsibility for different decisions. While scholars initially characterized "the profession" and "the bureaucracy" as different logics (Freidson, 2001) and "antithetical" (Freidson, 1984: 10), they have since recognized that the two divisions of labor are in fact compatible and increasingly seen structured together. Empirically, professions have been embedded in bureaucratic systems ever since the World War II-era when professionals started to become bureaucratized through salaried organizational employment and vertical integration (Barley and Tolbert, 1991). Today, most professionals are now employed in large bureaucratic organizations and subject to bureaucratic decision structures (Briscoe, 2007; Noordegraaf, 2011; Huising, 2015).

In this paper, we argue that many changes are coming about precisely because professional expertise is currently structured by and subject to those quantified bureaucratic decision structures. We found these to be in tension with algorithmic approaches being developed in data and technology-focused organizations. Our findings and argument are based on a 10-month ethnography of a retail technology company. Our study revealed a contrast between two conflicting ways that the company structured the expert work practices of a large professional group, with one way coming to replace the other. The original structuring approach involved using *bureaucratic decision structuring* to coordinate the collective decision-making of a more-than-200-person department of fashion buyers. The second structuring approach involved using *algorithmic decision structuring* to coordinate their collective decision-making.

Our data illustrate the tension and contrast between bureaucratic and algorithmic decision structuring. The bureaucratic decision structuring, organized by the organizational chart, structured decisions into roles and hierarchies as a single decision tree with unexamined impact that was not expected to change very much or very often. In contrast, the algorithmic decision structuring dynamically modeled and identified many ways of structuring decisions that were often at odds with existing roles and hierarchical structures and were instead based on likely performance impact. To resolve this tension and make use of the algorithmic tools, the buyers and planners had to reconfigure the role relationships and hierarchies structured by the organizational chart and disentangle them from the data structures. We propose that this fundamental difference between bureaucratic and algorithmic decision structuring offers a yet undertheorized reason why algorithms are bringing about large-scale structural change – especially because organizational charts and bureaucratic decision structuring have been prevalent modes of production for many decades in many industries.

THEORY

The structuring of professional expertise

Scholars have tracked the structuring of professional expertise through three key phases over time. First, Abbott and others recognized the creation and credentialing paradigmatic standalone professionals as a notable structuring of professional expertise. These practices by members of the profession created powerful solo practices as prototypical forms of professional employment (Cogan 1953; Goode 1957; Parsons 1968). Within those groups, members could exercise expert decisions with only the oversight of other members of the profession.

In the 1960s, as occupational groups that were considered “marginal professions” started to intensify their efforts to become “full-fledged professions”, scholars began to turn their focus to professionalism (e.g., Hall, 1968). A key focus of this research on professionalism involved how professionals establish divisions of labor by making claims to abstract knowledge and negotiating those claims (Abbott, 1988). Beginning in the 1970s, scholars became interested in how the power of the profession could explain the structural changes associated with the rise of professionalism (Berlant 1975; Parry and Parry 1976; Larson 1977). To explain these changes, these scholars drew on “market control theory”, which argued that professions sought market monopoly and social closure in order to increase their status in society.

The structuring of professional expertise in bureaucratic organizations

Following this work, scholars also began to recognize that the standalone paradigmatic professions that Abbott described were not characteristic of many ways that professional expertise was structured. This second wave of research recognized that, in fact, many professionals were employed by bureaucratic organizations, which meant that their expert decision-making was actually shaped and influenced by bureaucratic structuring, rather than just the self-monitoring and credentialing practices that Abbott described. Whereas scholars who had studied professional divisions of labor had examined expertise and effectiveness as central logics, those who had studied the bureaucratic division of labor acknowledged rationalization and efficiency as its central logics (Bunderson, Lofstrom, and Van de Ven, 2000). Over time, scholars have concluded that these different logics have implications for organizational design. Whereas research focused on professionals has tended to center on professional identity, standards, expert authority, and self-defined membership, research focused on the bureaucratic division of labor has tended to focus on specialization, hierarchy, and rules.

While researchers initially characterized the professional and bureaucratic divisions of labor as “antithetical” (Freidson, 1984: 10) and “necessarily conflictual” (Abernethy & Stoelwinder, 1994: 4), they later recognized that the forms may, in fact, be “symbiotic” (Hall

1968; Davies, 1983; Barley & Tolbert 1991; Montgomery, 1997; Bourgeault, Hirschkorn, and Sainsaulieu, 2011: 67). Abbott (1989: 279), for example, recognized that professionals and divisions of labor "can be nested inside one another." Weber foreshadowed this nesting by acknowledging expertise as a key premise of bureaucratic decision making, as well as a key determinant of organizational structure through "factoring" jobs into specific functions and roles that required high levels of expertise in specialized areas of specific organizational tasks (Weber, 1958).

Especially as scholars devoted attention to de-professionalization and proletarianization, they started to examine the implications of professionals being increasingly structured in bureaucratic organizations (e.g., Wilensky, 1964; Hall, 1968; Barley and Tolbert, 1991). Professions first gained prominence in bureaucratic organizations during the World War II-era when professionals such as physicians and lawyers, long considered prototypical professionals, transitioned from solo practitioners to working in bureaucratic organizations. From 1931 to 1980, self-employment among American physicians and lawyers fell from 80% to approximately 50%, and from 50% to less than 33%, respectively (Barley and Tolbert, 1991). Ultimately, multinational firms created legal departments, as well corporate medical units that were exclusively responsible for attending to their employees. Most professionals are now employed in large bureaucratic organizations (Hinings, 2005; Briscoe, 2007; Noordegraaf, 2011).

Professional expertise as control metrics infuse in bureaucratic structures

A third phase in the evolving literature on professions and the structuring of professional expertise relates to the fact that bureaucratic organizations have become increasingly characterized by quantified metrics of control (Fligstein, 1990). Such control metrics are now commonly used to structure, monitor, and evaluate the decisions of professionals. Fligstein

(1990) called this “conceptions of control”, defined as managerial paradigms of how best to solve competitive problems. Fligstein argued that, as social, economic, and regulatory environments in the U.S. have changed over time, different groups of professionals have sought to establish “stable” divisions of labor that enabled them to gain control of corporations and also exert control on their internal and external environments by establishing “legitimate action” (Fligstein, 1990: 6). According to Fligstein, through their “conceptions of control”, different professional groups have infused the modern bureaucratic organization with means of controls that have impacted specific aspects of organizational design.

Fligstein asserted that four different conceptions of control have emerged over time. First, in the years leading up to the twentieth century, the dominant conception of control involved direct control of competition, which was realized through the establishment of monopolies or cartels. Second, between 1900 and 1920, the “manufacturing conception of control” dominated as industrial engineer professionals entered factories, and managers and engineers gained control of firms. Trained in statistical and scientific methods focused on efficiency, these professionals introduced changes to organizational structures through the establishment of hierarchies and vertical integration. Third, during the Depression years, sales and marketing professions became the top executives in their firms as the “sales and marketing conception of control” dominated. During this time, firms shifted focus from price stability to product differentiation, and sought to effectively harness sales and marketing expertise by developing merchandising functions to coordinate production and organizing into divisions to enable the production of full product lines rather than narrow ones that had been characteristic of the manufacturing conception of control.

Most recently, according to Fligstein, the “financial conception of control” has reigned dominant as financial professionals, who are adept at managing firms as investment portfolios and evaluating the profitability of different product lines, have taken control of the firm. Fligstein demonstrated that, as part of the emergence of the financial conception of control, divisions were tasked with reporting their own financial metrics and implemented financial controls and advanced accounting systems. Today, most large organizations are dominated by the finance conception of control and "emphasize control through the use of financial tools" (pg. x). As Fligstein (1990:15) explained, "product lines are evaluated on their short-run profitability and important management decisions are based on the potential profitability of each line".

Bureaucratic rationalization reflected in the org structure

Taken together, the three streams of research reviewed above help us understand how professional expertise has been structured in bureaucratic organizations. Early scholars showed that, through the establishment of professional divisions of labor and jurisdiction boundaries, professionals' decision-making has been constrained to specific areas of expertise. Fligstein (1990), Barley and Tolbert (1991), and others showed that, as professionals have entered bureaucratic organizations, their expertise and decision-making structures have been subject to vertical integration, as well as bureaucratic controls. Dividing up tasks and decisions in this way is interpretable to humans. For example, the products that consumers purchase are produced by individual product lines that are structured in divisions in bureaucratic organizations.

ALGORITHMS AND PROFESSIONAL EXPERTISE

A separate but related stream of research has explored how data and algorithms have transformed the nature of professional expertise (e.g., MacKenzie, 2014; Christin, 2018; Lebovitz et al., 2020). Yet, to date, many of these studies explore professionals' expertise, but have not yet specifically examined how data and algorithms are changing the *structuring* of these professions such as their role structures or organizational hierarchies or networks (Barley 2020). Theories of technology and change argue hold that “revolutionary” technologies change how established roles and occupations relate to each other, suggesting that more needs to be understood about how these new data technologies are changing the structuring of professional expertise (Barley, 1990; Barley, 2015; Barley, 2020).

As our review of the literature above suggests, bureaucratic and organizational structures and financial controls have become primary to how professional expertise is organized in society. In this paper, we propose that these bureaucratic structures and related control metrics may be salient to how data and algorithms change the structuring of professional expertise in organizations. Even though many studies predict that data technologies will occasion significant organizational and societal change (e.g.), few studies have predicted or focused on how the Big Data revolution might affect these kinds of formal organizational structures. In fact, in 1978, Meyer explicitly said that automation—the use of computers as a technical innovation—would likely have little impact on formal bureaucratic structures. To date, few studies focused on data and automation have departed from that assumption. However, in the present study, we find such a departure. We studied the development of a new algorithm that automated work in a retail company and found that it unintendedly, but directly, impacted the formal bureaucratic structure. We analyze our findings to show how algorithms are now impacting formal bureaucratic structures, as opposed to strictly the jobs and processes that unfold within those structures.

Our findings stem from the observation that the decision structuring of algorithms is fundamentally different than that of bureaucratic systems. We define decision structuring as the process wherein individuals’ decisions are shaped by 1) decision rules, which are practices that prescribe what information to use when making a particular type of judgment, and 2) decision premises, which are prescriptions for how to reason about problems (Cyert and March, 1963; Simon, 1991). On one hand, bureaucratic decision structuring, as depicted in org charts, involves dividing out responsibility for different decisions into different jobs that are nested into managerial hierarchies. Bureaucratic decision structuring is driven by decision rules such as standard operating principles and decision premises such as roles that prescribe a unitary way to reason about problems. In contrast, algorithmic decision structuring involves decision rules and premises that produce many ways to reason about problems. Importantly, algorithms not only have the capability to newly measure and analyze the impact of the specific one way that the organizational chart is segmenting decisions, but they can also identify optimal ways of structuring the decisions, even as conditions change. Our findings, thus, illustrate the contrast between bureaucratic decision structuring which establishes a single stable decision tree with algorithmic decision structuring, which dynamically models and identifies many ways of structuring decisions based on likely performance impact. We propose that this contrast between bureaucratic and algorithmic decision structuring offers a yet undertheorized reason why algorithms are bringing about large-scale structural organizational change.

METHODS

Research setting: AlgoCo

The research site for this study was a large "digitally-native" retail technology company, to which we give the pseudonym AlgoCo. AlgoCo was founded in the early 2010s and, like many digitally-native retail companies, believed that "hyper-personalization" was the "new norm" in retail (Van Ossel, 2019). AlgoCo, however, differed from many online retail companies in its commitment to using proprietary algorithms for a wide range of applications, including predicting purchase behavior, forecasting demand, designing new apparel, generating marketing strategies, and optimizing inventory. At the time of our study, AlgoCo employed over 100 data scientists in a centralized, powerful "algorithms" department. The algorithms department had developed and implemented more than 100 algorithmic capabilities that were "in production" meaning recommending or automating work and decisions throughout the organization. Since its founding, AlgoCo had focused significant resources on developing a proprietary data infrastructure that included detailed information on distinct units for sale, or "Stock Keeping Units" (SKUs). In addition to incorporating basic SKU data such as clothing material and size, the data infrastructure also encompassed details on client preferences, including whether SKUs aligned with the preferences of each individual consumer, as well as client feedback, including ratings of how well a certain piece of clothing fit. One data metric that was important in this study was fine-grained data on customer-item outcomes, for example, the "keep rate" of each item, which calculated the number of times that a customer kept (versus returned) an item divided by the number of times a customer was sent an item. This metric was fine-grained, in the sense that it calculated the keep rate of items by different colors and different sizes. So, for example, anyone could look on a main dashboard to see whether that the navy blue size XL version of a sweater had a better keep rate than the green size S version of that same sweater.

We selected AlgoCo as the context for our study because it was well suited to developing new theory about how organizational structures are impacted when a company relies extensively on algorithmic capabilities. Our observations at AlgoCo began in 2017. At that time, AlgoCo had almost 2000 employees, of which about 100 people comprised its "algorithms" department. AlgoCo's commitment to its proprietary algorithms was reflected in its organizational chart, wherein the algorithm department was centralized rather than "folded" into other departments, as is common among many other retailers. Unlike most retailers where data science and algorithm departments—if they exist—work in service of other departments, at AlgoCo, the algorithms department was managed by a Chief Algorithm Officer who, in turn, reported directly to the CEO. The algorithms department worked autonomously, with each data scientist in the department "owning" his or her own research project focused on developing transformational algorithms. They owned the algorithms even as they were "put in production" meaning part of the live set of platforms that the employees throughout AlgoCo used to do their work.

The Assortment Planning algorithm

The data science team that we studied included two experienced data scientists and a user interface (UI) designer, and later grew to include two additional data scientists. Their work involved developing new ways of mathematically modeling the assortment planning process – which is the way that fashion buyers plan out the inventory that the company should develop, produce, or purchase from vendors to then sell to customers. Even though the team we studied was specifically focused on developing a new assortment planning algorithm, when they explained their work to others, they often described how their algorithm interfaced with a whole interconnected algorithm system that drew on centralized data (see Figure 1). The team was trying to formalize the assortment planning process as a constrained optimization problem, which would mean that their model could include and weigh all of the decisions about all items of clothing in the same mathematical formula (which would, in turn, be embedded software code) that could calculate how different decisions impacted the overall optimization score of a set of inventory decisions. In parallel, while the data scientists were developing this formalization of the problem, they were also meeting weekly with the UI designer who was creating new user interfaces that the buyers could use to visualize this new formalization of the assortment planning process. As we will elaborate in the findings, we did not see a lot of resistance to using the tool based on occupational identity or autonomy issues as has been found in prior studies (Christin, 2018; Kellogg, Valentine, & Christin, 2019), but we did observe consistent mismatches between the algorithmic and bureaucratic approaches to understanding and structuring decisions.

Insert Figure 1 here

Ethnographic data and analysis

We used an inductive, ethnographic research approach in this study. Our research design was guided by the open-ended question of whether and how newly developed algorithms were changing professionals' work. Because this area of research is new and growing, an inductive field-based approach was well-suited to the question (Edmondson and McManus, 2007). The first author negotiated access to work as an unpaid program manager for the AlgoCo algorithms department. The unpaid program manager position allowed access to company headquarters in a large US city, and to all of the company's internal data, communication, and coordination platforms. It also involved going through the onboarding and socialization processes, and joining the algorithms department community, including weekly happy hours, as well as the company-wide community, including weekly company-wide meetings.

In addition to providing the program manager administrative services, the first author negotiated access with the one specific data science team that was developing one specific algorithmic capability. Embedding in their team involved attending local team meetings, and their managers' meetings and directors' meetings, which gave us a broad view of the work of many related data science teams. The algorithm that our data science team was developing, described in detail below, was focused on a particular workflow (assortment planning, see Figure

2) within the job of fashion buyers at AlgoCo. Fashion buyers were paired with planners who did much of the computation in support of the buyers' job. Therefore, the first author also identified a specific fashion buyer-planner pair who agreed to let us study their work before, during, and after the algorithm was developed and deployed. This shadowing included following the fashion buyer to a New York City buying trip where the fashion buyer met with vendors and negotiated prices and orders with many vendors. The fashion buyer-planner pair was embedded in a larger buying team, which allowed us to collect data from all of their weekly team meetings and department meetings throughout the 10-month period that we observed. This focus allowed for a clear analysis of the workflow before, during, and after the algorithm was developed.

Insert Figure 2 here

The data science team took a "human-centered design" approach to the tool development, first observing the buyers' work, and then engaging collaboratively with them to understand their needs and mental models of their assortment planning work (e.g., Cooley, 1999; Norman, 2005). This collaborative approach meant that we could also analyze extensive data on the cross-functional interactions between the data scientist local team and fashion buyer pair, as well as between the data science managers and buying directors. Many of these interactions took place either at user testing meetings, cross-functional governance meetings for this specific algorithm, or later user training meetings. Finally, as we began to understand that a key finding related to the tension between the algorithmic and bureaucratic decision structuring, we also collected archival data on the fashion buyers' org chart or department structure over the decade that AlgoCo had been in business.

Our main argument relates to the significance of changes that we observed over time.

Our analytical approach was therefore structured to characterize, substantiate, and illustrate changes over time. We conducted a thorough analysis of all our observations of the buyers' and planners' work as it was enacted using all of the AlgoCo data platforms before the new algorithm was developed. We analyzed how decisions and metrics were used, discussed, and interpreted during daily work and regular meetings. We also analyzed how change was understood and accomplished within this earlier phase of how the buyers worked, because those within-phase changes reveal many assumptions about why things "were the way they were". Within this first phase, before the algorithm was developed, people had a taken-for-granted way of making sense of their decisions, jobs, metrics, and many of those taken-for-granted assumptions were surfaced in discussions during minor changes to decision structuring within this phase. We also conducted a thorough analysis of the many cross-functional interactions that played out as the data science team developed the algorithm, in collaboration with the fashion buyers. These interactions began to surface many of the tensions that are the focus of our paper. We analyzed the discussions, tensions, and resolutions that played out in each meeting and interaction during this period. Finally, the first author also returned to AlgoCo a year after leaving the field to follow up on the adoption and use of the algorithm, and to see how some of the tensions had evolved and resolved. The final phase included analysis of all of the interviews, meetings, and observations conducted during the month of follow-up with the data science team and their managers, and the fashion buyers' teams and their managers.

FINDINGS

This section contrasts two different ways that AlgoCo structured the expert work practices of a large professional group, retail fashion buyers, with one way aiming to replace the other. The original structuring approach involved using *bureaucratic decision structuring* to coordinate collective decision-making among the large group of professionals. In this system, each professional occupied a role that had authority and accountability for a codified and specific set of decisions, and then groups of such roles were nested in managerial hierarchies. The managers had authority and accountability for the collective sets of decisions made by that group of roles, as is typical in bureaucratic organizations. The organizational chart not only structured jobs and hierarchies, but also financial and performance metrics, which were assigned to each role and each manager and thus followed the structure of the org chart. Each role used their metrics to guide their work practices on an ongoing basis, such that the metrics enabled the large group to carry out coordinated decisions.

The second way of structuring the expert work practices was different and aimed to replace this original way. The second structuring approach involved using *algorithmic decision structuring* to coordinate collective decision-making among the large group of professionals. A newly developed algorithmic tool modeled 'scenarios' – meaning it computed the likely performance outcome of different decisions. At first, the algorithmic tool was used for algorithmic decision structuring within each role; each professional could use it to optimize the decisions they made for their own local metrics. However, this exercise began to show that the way each role's decisions and metrics were divided out also influenced how the optimal outcomes were identified. This realization then began to call into question the way decisions and metrics were structured among all of the roles and hierarchies – the structuring of the overall organizational chart itself. The algorithmic tool newly revealed the performance impact of different ways of dividing out and nesting the metrics among the different roles and hierarchies, with two implications: first, the formal bureaucratic decision structuring process itself became subject to new measurement and analysis. And second, the algorithmic tool not only demonstrated the impact of the *unitary* way that the organizational chart was dividing out the metrics, but it also modeled "arbitrarily many" ways of structuring the decisions and metrics, helping identify many optimized ways of structuring the decisions – prompting a move towards more dynamic, flexible, and responsive structuring. These findings reveal the contrast between bureaucratic decision structuring – which sets up a single decision tree with unexamined impact that is not expected to change very much or very often – and algorithmic decision structuring which dynamically models and identifies many ways of structuring decisions based on likely performance impact (see Table 1).

Insert Table 1 here

The rest of the findings more fully illustrate this contrast and the differences and tensions between the two ways of structuring decisions. We first use our ethnographic data to explain how a more-than-200-person merchandising department used bureaucratic decision structuring to coordinate and accomplish "assortment planning," a complex and collective decision-making process. We then describe how data scientists developed the new algorithmic tool to support this process, which began to surface the contrast and tension between bureaucratic and algorithmic decision structuring. We conclude the findings by describing how data science and merchant leaders worked to understand, disentangle, and make use of both the bureaucratic and algorithmic decision structuring processes for assortment planning and related work.

Phase 1:
Professionals Use Bureaucratic Decision Structuring to Coordinate Collective Work

The merchandising department's organizational chart and related metrics structured the decisions domains and work practices of about 200 merchants in various 'buyer' and 'planner' roles and in different levels of hierarchy. Figure 3 depicts a stylized version of the org chart.

Insert Figure 3 here

Assortment planning process. There were many related processes involved in developing and maintaining a well-performing inventory. Here we describe the specific assortment planning process under bureaucratic decision structuring. About six months before a fashion season began, the planning, strategy, and marketing executives determined a set number of buys for all of the merchandising department and set the financial and performance metrics for the entire group. We can narrate this process with falsified and simplified numbers, as illustrated in Figure 4. The buying executive might be given 600,000 buys and a keep rate (KR) target of 75%. The work of selecting and developing the 600,000 buys was complex and, so, it was further divided such that each buying director was allocated 100,000 buys each with KR targets level-loaded across departments based on the expected performance of each department. This process of dividing out the buys and level-loading the metrics continued to the "bottom" of the organizational chart with the 100,000 units divided up among the front-line buyers who negotiated with vendors. In our simplified example, the Plus buying and planning managers decided to allocate 30,000 units to Wovens, 20,000 units to Knits, 10,000 units to Dresses, 20,000 units to Denim, and 20,000 units to Bottoms. This complex process was data-driven and dynamic. Much of this work was accomplished using Excel (or Google) spreadsheets that had been programmed with sophisticated macros (automated input sequences that calculate complex formulas across different cells and tabs in a spreadsheet) that helped calculate the potential impact of moving a set of buys from one product category (e.g., men's denim) to another category (e.g., women's knits).

Insert Figure 4 here

Buyer-planner workflow in assortment planning. The assortment planning process thus involved a "top-down" allocation of buys and targets that defined the work and targets of fashion buyers at every level of the merchandising org chart. Continuing our stylized example from above, this process meant that, at the beginning of the assortment planning process, the buyer-planner pair who was in charge of the women's plus-sized denim category would be given 20,000 units of denim "buys" for that quarter and a set of performance targets that they needed to hit with those buys. They followed a somewhat similar process as the overall buying and planning directors as they developed an assortment based on their 20,000 units—they all relied on the similar spreadsheets that were pre-coded with macros, and they dynamically and iteratively moved their buys across different product categories to see what the predicted impact would be on performance targets. At the director level, this process meant seeing the predicted impact of moving larger sets of buys between the women's versus men's department. At the front-line buyer-planner level, this process meant seeing the predicted impact of moving sets of buys between different kinds of plus-sized denim (e.g., capris versus boot cut) and between different vendors.

Figure 5 offers a stylized illustration of what this local assortment process involved. The buyer-planner pair met and discussed frequently how to allocate the upcoming seasons buys across many different denim product categories, including style (capri, skinny, boot cut, straight), color (white, light, medium, dark, stonewash), price point (under $50, $50-$80, $80+), and vendor. Each of these product categories had different historical performance, which was the source of data used to predict assortment plan performance. The role relationship involved the buyer being more of an artist with an intuition for upcoming trends, and the planner making the vision work by moving buys around between the different product categories.

Insert Figure 5 here

Organizational changes under bureaucratic decision structuring. During our data collection, we observed or learned about three instances of organizational change related to the buyers' organizational chart. These instances provide a useful analytical lens for understanding how the organizational chart was being used and its assumed purpose. The discussions around these changes illustrate how the "org chart" and related bureaucratic decision structuring was dividing up the large set of complex decisions and related tasks to be manageable and interpretable for humans.

The first example involved creating a new role on the Plus buying team as the volume of purchases in that customer segment grew. Originally, the Plus buying team had a buyer-planner pair who planned and managed the assortment for "Tops." As Plus sales volume grew, it became infeasible for one buyer and planner team to make all the purchasing decisions for that category and so the "Tops" category was split into two subcategories—"wovens" and "knits and sweaters". The Plus Buying Director explained the decision,

> We split out Tops into someone who was responsible for wovens and someone who has responsibility for knits and sweaters, just to make the scope of responsibility more equitable and more manageable.

The decisions and targets for the Tops buyer were thus segmented into two buyer roles—one buyer-planner team was responsible for developing the inventory for wovens, while another team was responsible for knits and sweaters. Each buyer-planner pair was assigned their own volume and targets. If the Plus Buying Team as a whole was allocated 100,000 units, the wovens buyer might be allocated 20,000 units and the knits buyer might be allocated 20,000 units of her own. There was no discussion of whether this division would impact the outcomes or targets, it was an assumed, taken-for-granted division of labor based on the growing sales and need to split the number of decisions for manageability.

The second example occurred when we were conducting our observations and involved the company newly entering a new market in the European Union (EU). The merchandising department expanded to include an EU department alongside the women's, men's, plus, and kid's departments. The EU executives who formed and structured the department decided to structure the buying teams and roles based on how the customers might use the clothes, rather than on the more standard product type. The EU buying team thus had an Evening Wear buyer, a Casual Wear buyer, and a Workwear buyer. A buyer might include a dress in developing an inventory for each of these categories because dresses might be appropriate for any of these uses. In contrast, the US Women's buying team had dedicated Dress buyers who would purchase all of the evening wear, casual wear, and workwear dresses. As the EU buying department was being structured, this non-standardized way of structuring the buying roles was easily accepted by the merchandising department and executives. It was explained to be the way of structuring and dividing out the decisions that was most manageable and useful for the EU buyers. Later, the non-standard roles and product categories introduced complications for some of the data science approaches, but under the original bureaucratic decision structuring work, this division was straightforward. There was no discussion of whether this way of dividing out the decisions would impact targets or outcomes.

The third example had happened a year before we began our field observation, but many people discussed it in interviews, and we collected extensive archival data on this example. This change involved a large "re-org" (an emic term, short for "re-organization") that changed the buying teams and reporting lines in the women's department. This example again illustrates how these bureaucratic structuring decisions were made based on the need for manageability of large areas, and were made to support interpretability and clear accountability for a decision domain. AlgoCo had been founded with only a Women's department, then later added Men's, Plus, Kids, and EU departments. This comparatively long history meant they had higher sales volume in this Women's category. The higher sales volume had become unmanageable – every buyer-planner pair was overwhelmed by the amount of inventory they had to develop and manage. So the executive who led the Women's department felt it was time to divide the department into smaller areas so that teams could more easily manage those smaller areas.

One merchandise executive reflected on a strategic opportunity to better use the data insights as she thought through this re-org:

> We were sub-optimizing the buyers' decisions because they were gravitating toward the average. But the average of a big base…That was not serving our clients, particularly those at the bookends of the spectrum, whether it's age, or price preference, or style.

Her sense was that defining a buyer role by product type meant they bought those products with the average customer in mind. There were many conversations about how to use the re-org to pursue this strategic opportunity of having the buyers focus more on specific customers. For example, at a multi-day "off-site", executives, merchant leaders, and data scientists all discussed how to split up the Women's department.-The data scientists in the algorithm department wanted to divide up the department by customer age segments so that the buyers could focus on developing inventory specifically for different age groups. They defended this proposal by arguing that age was the client attribute that most significantly predicted KR. They argued for structuring based on which segmentation related to client outcomes, not based on how buyers and planners think about or interpret their work. One data scientist explained,

> I wanted to buy by age segment to introduce a source of diversity into our assortment...I focused on Age because another Algos team had shown that Age was the client attribute

that most strongly conditioned Keep Rate.

In contrast, the merchandising team wanted to divide up the women's department based on price point. They thought that focusing buyers on developing inventory within "low price point denim" would be a better approach for dividing up the department, and also for introducing more diverse and targeted inventory. The merchandising teams' reasoning was based on intuition. One of their executives explained that she did not think that customers' preferences were that different based on their ages, so she thought developing inventory targeted to the ages would not produce better inventory. In the end, the merchants' authority for their own department and workflows prevailed, and the Women's department was divided up into bargain, general, and luxury price points. The buying directors of these sub-departments reported to the Women's exec, and the "buys" and targets were divided out among these newly formed buying groups.

Figure 3 which shows a stylized version of AlgoCo's "org chart" illustrates that the Women's department has one more level of hierarchy than the other departments. This re-org was responsible for that additional layer in the Women's department hierarchy. The overall org chart of the full merchandising department involved non-standard categories and levels for dividing up the decisions. But it provided relatively clear accountability for decision domains within the large, complex, coordinated set of decisions involved in planning and managing a massive inventory for a large and diverse set of customers.

Phase 2:
Tensions between Bureaucratic and Algorithmic Decision Structuring

We next describe the change process that unfolded as a data science team developed a new algorithmic tool for the buyers to use during assortment planning. We observed a consistent interaction pattern during this process that was especially evident during cross-functional meetings such as user testing meetings, the weekly planning meetings, and later at the user trainings. Over time, and through our analysis, we realized that this interaction pattern related to fundamentally different ways that the buyers and data scientists were approaching the decision structuring process. To explain this contrast, we will illustrate how this tension emerged within individual buyer roles and then within the buyer hierarchy.

Algorithmic decision structuring in tension with the professionals' roles

The tension that arose related to the buyers' individual roles was related to their use of non-standard and practical product categories to guide their decision-making. The categories were non-standard in the sense that they were not the same across the different buyers' roles. That the categories were non-standard across buyers was not salient under bureaucratic decision structuring, because the buyers were responsible for "hitting" higher-level targets so that the focus was on hitting those metrics rather than the ways in which they went about this, which were not attended to. In contrast, the data scientists were expecting more standard categories to be built into the tool, and expected that the categories that would be preferred would be the ones that predicted outcomes. The data scientists also expected that client categories (for example, age segments), not just product categories (for example, denim styles) would shape decision-making. These tensions were apparent at the first user testing session and shaped most interactions during the tool development process.

How to structure decisions. The first user prototyping session involved the buyer for Men's bottoms. Note that the Men's department had not achieved high enough sales volume to warrant a dedicated Men's denim buyer as in the case of the Women's bargain and general departments, so this buyer covered and was responsible for all decision making related to Men's bottoms, including denim. The user testing session involved her inputting her buys into the tool and getting the algorithmic recommendation for how she should allocate the buys. After she ran the tool, she immediately noticed that the tool was splitting up recommendations in a way that she did not typically do. Looking at the tool, she reflected,

> Today how we buy is, for every style, we buy across all inseams. So this set of recommendations is super different from what we'd actually do today. Rarely do we buy one pant in one inseam. This recommendation is telling us that certain styles would be better in certain inseams.

The data scientists' initial response was representative of their approach. They first explored whether this different way of splitting up the decisions (i.e., buying different styles in different inseams rather than buying all inseams for any one style) might be predictive of outcomes. They did not feel as constrained by the practical product categories the buyer was using, and were expecting that something in the data was guiding this algorithmic recommendation. One of the data scientists suggested that it was possible that on the client side that shorter clients (i.e., smaller inseam) preferred a certain color and style of pants and taller clients (i.e., longer inseams). The buyer admitted, "It's a valid question" acknowledging that this other set of client categories (tall or short clients) might explain their buying patterns and therefore the recommendation she was seeing. She then explained the practical reasons why she needed to make the decisions in a way that did not split up buys across inseams and styles, including to negotiate with vendors. Together, she and the data scientists decided that later they might be able to investigate the question of vendor negotiations and vendor minimums, but for now, they would try to support the buyers' existing practice of buying all inseams across all styles; this required adjusting both the model and the interface of the tool.

As they explored the issue, the data scientists also were learning that this issue of non-standard and practical product categories would likely arise for every buyer. One of the data scientists tried to discover how generalizable this inseam issue would be. He said, "Got it. So (the tool is separating out styles and inseams) and that is a big blocker… Hmm. That's for men's bottoms but not women's…?" The buyer replied, "Yeah, women's doesn't use inseams."

We can connect this example to the main tension. In the past, it did not matter whether she separated out her buys by inseams and styles or whether she and the women's denim and pants buyers did that differently. What mattered was whether their part of the inventory was profitable and performing well with clients. When the algorithmic tool recommended splitting up inseams by styles, the data scientists guessed it was because of an underlying pattern in the data where different clients with different heights preferred different styles and wanted to try buying in a way that reflected that data pattern. The buyer could not accommodate that recommendation because it did not fit with her practical reasons for buying all inseams in any one style. This men's bottom was the first user testing session, but each subsequent session revealed additional practical and non-standard categories that were important in how the buyers planned their assortments. Even the last user testing sessions that we observed in the last month of our observation, which was for Plus dresses, involved new product categories. That buyer wanted to be able to structure her assortment decisions across many dress silhouettes that had not previously been raised by any other buyer.

Depth recommendations. This example relates to the broader tension that arose during this phase of tool development, which related to differences in what the buyers assumed was a good assortment versus what the algorithmic tool recommended as a good assortment. A specific tension was around the "depth" of styles that the buyers assumed made up a good assortment versus what the algorithm recommended. Each buyer was given buys and targets and had to distribute those buys across a variety of styles. Each buyer used practical, intuitive—and non-standard—product categories to divide out across their chosen number of styles (e.g., bootcut, straight, skinny, capri) and to the depth of buys in any one style to create what they considered a good assortment. The algorithmic recommendations began to explicitly model and measure the impact of their intuition of a good assortment.

The Plus wovens buyer's first user testing session illustrates the tension. She and the data science team had this exchange:

> **Buyer:** …Because when I ran it the first time it basically told me I should buy every single unit in this first style which, truthfully, isn't even that good.
>
> **Data scientist**: Oh that's interesting. Which one was it that they've put too much into?
>
> **Buyer:** (gesturing to screen) If I didn't put this 1000 unit threshold, what I was getting was… it would say put 10,000 units into one style and put 10,000 into this other style.
>
> **Data scientist**: Right, so that is, again, the expected behavior because that's ... if you don't tell the tool to force some amount of breadth the tool is just going to be like, put everything in the best one.

The data scientist was explaining that an optimization algorithm was going to optimize in a simple or naïve way— it was going to find the item with the highest keep rate and recommend that the entire set of buys be allocated there, because it would indeed optimize the keep rate. The buyers' concern was that putting 10,000 units to that one style that had performed well in the past would not represent a good inventory. There were many conversations among the buyers outside of the user testing sessions about how comical it was that the tool was telling them to buy 10,000 units of one style. The buyers and planners discussed these initial depth recommendations as a big failure on the part of the algorithm.

The data scientists were agnostic as to whether 10,000 units of one style was a good inventory or not; to them, that was a question that would be best answered by data. The ongoing tension between the two approaches was not whether 10,000 units was the right number of units, but the higher-level question of how to understand what a good assortment was. The data scientists' approach was to use historical data to introduce diversity into the assortment—they wanted the optimization algorithm to be constrained by empirical data patterns such as customer segments, seasonality, and vendor relationships. The buyers' approach was to use their intuition and experience to introduce diversity into the assortment—they wanted the optimization algorithm to be constrained by their intuitive product categories such as what styles, colors, silhouettes should be included each season's assortment. The tension was not directly interpersonally conflictual. Instead, it involved very different approaches to thinking about assortment planning.

Product categories. Another example of this tension involved figuring out whether and how customer categories should be used in the assortment planning process. As explained in the first findings section, although buyers consulted customer data as they were planning, they did not do so in a holistic or systematic way. Instead, they would look at the customer data pertaining to any style that was performing in an unexpected way (recall the "under the covers" phrase). The data scientists wanted to include customer categories in the constrained optimization model for assortment planning. They thought that recommendations would be more diverse in a way that reflected customer preferences if the assortment decisions were constrained and optimized for customer segments. The data scientists wanted to include age-related client segments in the calculation of the algorithmic recommendations. The following example illustrates the tension. In a follow-up user testing session, the data science team introduced the fact that they were including client segments by age in the recommendations now. This conversation illustrates:

> **Data scientist 1:** One slight change... I did add client segment into it. We do not have to include that. Last time you said it does not change that much. *Including it allows us to see if that's true.* If we don't have client segment in it, then KR will be biased to be higher because it will take the 50 and over segment.
>
> **Data scientist 2:** Because it's maximizing KR and the older age segment has higher KR.

> **Buyer:** I guess that's important to know because right now we don't sort too much by age group. The [bottoms styles that we buy] that do well… do well across age group. But knowing that would impact the end result is important to note for whoever is using it.

Her response was indicating her assumption that the way to use that information would be for the stylists who decided what styles to send to what customers to know that a particular style might perform well with different age groups. She did not necessarily agree that her assortment planning process should include considerations of age. She continued, "I wonder if we even don't have metric this in…" meaning that the set of recommendations should not even be informed by historical buying patterns separated out by client age.

Age was not the only client variable that predicted variance in terms of purchase patterns, and the data scientists were interested in including other client segmentations in the assortment planning process as well. For example, another data science team had used machine learning algorithms across all customer and item interactions to analyze the latent factors that contributed to a customer's preference for an item. That team had determined four different clusters of clients that determined variation in purchase patterns, including client styles such as "Casual" or and "Classic". The assortment planning data science team suspected that assorting based on latent style (versus for the average customer, which is what they thought the baseline buying process was aimed at) might improve inventory performance.

However, like age-based client segments, this segmentation increased complexity on the part of buyers and planners and increased their workloads, even though it predicted more variance. Earlier attempts – years before – to get the buyers to plan based on customer segments had failed because the client segments had introduced too much complexity into the planning process. The buyers still wanted to use their original product category strategies ("a well-performing denim inventory includes skinny, straight, boot-cut, boyfriend") so including customer categories to them just meant three times the work for unclear benefit. Figure 6 illustrates the problem of trying to include client categories in the original process that used spreadsheets for assortment planning. In describing this problem at a meeting with a large group of merchants, a data scientist said, "Rebuys are really successful… when they are purchased for client segments, they are extra successful. And not just for KR, but for lots of other aesthetic and financial targets. But in the past, forcing buyers to think about client segments makes the (client segment insights) less scalable."

Insert Figure 6 here

Algorithmic decision structuring in tension with the professionals' hierarchy

In addition to the tensions related to the buyers' roles reported above, the second set of tensions related to how the multiple and complex decisions involved in assortment planning should be divided and aggregated. This related to the buyers' organizational hierarchy structured by the org chart. Again, the differing approaches were about practical, interpretable, and fairly static structuring on the buyers' side versus more flexible structuring that explicitly predicted outcomes on the data science side. For the data scientists, the structuring that would be preferred was always that which had predicted outcomes in the historical data.

The org chart is a decision tree, decisions made in 'leaf nodes'. A key idea to understanding this tension is to understand why the data scientists considered the merchandising organizational chart to be a decision tree. Several data scientists in various meetings talked in offhand ways (meaning most people there understood the point) about how the merchandising department data structure and organizational chart, which were structured around the same product taxonomy (recall Figure 3), could be understood as "decision trees." They explained that within the decision trees “the buying all happens at the leaf nodes.” One of the data scientists elaborated on this point in an interview. He showed us a data interface that organized all the items in the AlgoCo inventory. He showed earrings as an example product category,

> So there's the ID, there's the name, there's a parent ID. So, earrings has ID 63 and has a parent that is 9, follow that and. see earring's parent is jewelry which is ID 9. And jewelry has a parent whose ID is 87… accessories.
>
> See how earrings has a parent (jewelry) and jewelry has a parent (accessories) There are some things that if you follow down, *nothing has them as a parent. Those are leaf nodes.*

He then emphasized, “So those (gesturing to leaf node) are the groups that actually go out and buy things. And then the others are just roll-up groups.” He was referring to the fact that the buyers who made buying decisions were at the “level” of jewelry. Actual purchasing decisions were not made at the “roll up” levels like accessories. He explained further, gesturing to his screen, “There are people *here*” (gesturing to the buyers) that actually buy stuff. And there are people *here* (gesturing to another buyer in the same group) that buy stuff. But *here* (gesturing to their manager and their manager’s manager) there's no one here that buys stuff.” He concluded, “The budget for this leaf node (meaning the buyer) and this leaf node (the other buyer) roll up to the budget for this parent node (the manager). But no buying happens here (at the manager level).” One of the data scientists on the team we studied connected this idea to their algorithm:

> If you have a hierarchy where information flows bottoms up and tops down like this, where the decisions happen here, here, and here (indicating leaf nodes and the roll-up teams) rather than side to side, you are naturally going to have workflows that have to involve leaf nodes and these bottoms up decisions.

She further explained that other algorithmic design processes could look at “hooking in at other places where the information might be flowing. But for us designing for this buying decision meant designing at the leaf node.”

Org chart hierarchy was constraining optimization. The data scientists had chosen to design at the leaf nodes because that was where the buying decisions were made. However, an issue soon arose because it became clear that the structuring of the leaf nodes was somewhat arbitrary but was influencing what the algorithm could recommend. We can report a simple example to illustrate and then explore this insight and its implications more fully. To check our understanding of this dynamic, we asked one of the data scientists in an interview,

> **Interviewer**: OK so what you all are saying is… Consider two scenarios. In the first you set up two buyers' roles: 1) women's workwear and 2) women's casualwear and give them each 1,000 buys… and then run the optimization algorithm on the 1,000 within workwear and the 1,000 within casualwear.
>
> In the second scenario, you set up the two buyers' roles as 1) women's tops and 2) women's bottoms and give each of them 1,000 buys… and then run the optimization algorithm on the 1,000 within tops and the 1,000 within bottoms.
>
> You're saying that in these two scenarios, you would get a different set of recommendations… and you would stock a different inventory.
>
> **Data scientist**: It seems most certain that you would.
>
> **Interviewer**: And one way of doing it would produce better outcomes.
>
> **Data scientist**: Right. And you could measure it.

As this quote illustrates, the data scientists and buyers began to realize that the org chart itself was segmenting out decisions into jobs through the assigned buys and targets, and in ways that influenced what an optimal set of decisions would be for that job. The merchandising org chart was segmenting out decisions to be made in a way that, according to the data scientists' approach, was unexamined and not optimized. Instead, the org chart was structuring the individual buyers' jobs and related buys and targets based on an interpretable and taken-for-granted product taxonomy (Figure 3). The directors would use the interpretable product taxonomy to divide out targets, and the buyers would hit their targets. Most of the focus was on whether the buyers were hitting their targets, not where the targets came from and whether a different or better set of targets were possible. The algorithmic approach newly called into question the process of dividing up the targets because that defined the space of decisions on which the algorithm was optimizing. The algorithm would optimize within whatever set of decisions it was given, and the buys and targets—following out the interpretable product taxonomy of the org chart—were defining decision space that would be optimized.

This issue was not contentious between the two functions. Rather, both buying directors and data scientists recognized it as a problem and discussed it in meetings and interviews. As an example, a buying director said, “Segmenting our teams and inventory in this way doesn’t allow for our algorithms to explore scenarios about our inventory and our clients in a multidimensional way. It also does not let us optimize for multiple performance metrics.” Her point was that the algorithmic tool was set to explore the 60,000 buys of one buyer. But that same approach could be used to explore across 120,000 buys of two buyers, and so on, to see if there was an optimal pattern across their buys. A data scientist expressed a similar idea this way: “We saw the algorithm could explore a larger space for better results.” One of the executives said in a strategic planning meeting after they were discussing the merchandising org chart constraining the algorithmic exploration, “I fully understand the drawbacks of how we are currently organizing the merchandising department. We are just now figuring out the better way.” A final quote illustrates how this problem related to a core principle for the data scientists. Several of the data scientists had heard in their disciplinary training the phrase, “Binning is sinning” which referred to the idea that data should be modeled as a continuous distribution and that imposing “bins” or categories on the data would introduce a lot of distortions and problems. One of the data scientists suggested, “You’ve heard the phrase ‘binning is sinning?’ I wonder if this is an artificial binning... We might be at a temporary period in the history of AlgoCo in which we're artificially binning the way we are buying as opposed to buying for specific clients."

Phase 3:
Reconfiguring Roles and Hierarchies to Accommodate Algorithmic Decision Structuring

The section above illustrated the tensions that arose as the data scientists' approach to decision structuring interfaced with the buyers' bureaucratic approach to decision structuring. Their mutual attempts to accommodate algorithmic decision structuring involved reconfiguring role structures and organizational hierarchies – i.e., the kinds of major structural changes predicted by theories of technological change for more "revolutionary" technologies (Barley, 1990; Barley, 2015; Barley, 2020). At the end of the two-year study period, the algorithmic tool had been broadly adopted, some changes to the role structures and hierarchies had already been accomplished, and other changes had been specifically planned and discussed. These involved changes to the buyers' role, their role structure with the planners, and the hierarchy that had structured the merchandising department (see Table 2).

Insert Table 2 here

Reconfiguring the professionals' role and role structure

As described in Phase 2, the buyers and data scientists differed in how they thought about the relevant categories that should be used to diversify any one buyer's assortment plan. This tension resolved in how the data scientists designed the algorithmic tool, with implications for both the buyers' role and their relationships with planners.

The buyers' role. The first tension related to the buyers' use of non-standard and practical product categories to guide their decision-making, when the data scientists expected that the categories that would be preferred would be the ones that predicted outcomes. One resolution to this tension involved the data scientists configuring the algorithmic tool to let the buyers input any of the practical and non-standard categories that they wanted into the tool in the form of constraints on the optimized recommendations. This change involved including more data. As one data scientist explained, “We added whatever targets they asked for: brand, color, silhouette, price, print, fit, inseam, rise, latent style. Every buying group wanted different ones.” Accommodating the product labels meant that the buyers more easily transitioned to using the tool, and it also helped the data scientists collect data on the different product attributes that the buyers considered relevant to developing a well-performing inventory. This change then also included changing the tool slightly to allow the buyers to input any of those product categories as constraints on the recommendations. As an example, after the exchange reported above where the data scientist explained to the buyer that the algorithm would allocate all units to the best style if not forced to add some breadth of styles, she added, “And so that's like where you get to be creative – to pick out what kinds of breadth do you want”, encouraging the buyer to input her intuitive constraints to create her idea of what a good assortment would be. To support this desired behavior from the buyers, the data science team configured the buyers’ interface so that they could enter constraints such as Vendor A must supply 30% of inventory, the Bootcut style must comprise 70% of inventory, or red-colored Denim must comprise 10% of inventory. This functionality allowed the buyers to segment their assortment by any of these categories and visualize how various depth and breadth structures impacted their targets. Similar to the product category labels, this functionality was seen as fairly straightforward for the data scientists to add to the tool. One of the data scientists explained this in a presentation at a cross-functional meeting. She said, “We’ve gotten a lot of feedback along the way. The financial targets and aesthetic and breath targets were easy to implement. We are all still learning how to combine these to deliver assortments that match people’s intuition.”

Note one key implication of this reconfiguration of the tool. The feature that allowed the buyers to input their constraints in order to structure the assortment the way that they wanted meant that their intuitive structuring of their assortment became explicitly modeled and measured. If they had an intuitive sense that structuring their assortment to include 30% of their units in a certain category—say bootcut denim—that structuring decision was now recorded and measured. In the past, they were accountable for whether they hit their targets, but with little understanding of how they were structuring their assortments to hit their targets. Now, their intuition pertaining to how to hit their targets was recorded and could be measured and analyzed over time. Even though the tool allowed them to plan inventories based on practical, intuitive categories, it also implicitly subjected those to the data science approach of assessing the impact of that decision structuring on outcomes.

Recall that a second related tension was that the data scientists also expected that client categories (for example, age segments), not just product categories (for example, denim styles) would shape decision-making. The resolution to this tension involved the data scientists including client segments "behind the scenes", meaning client segments were added to the code that produced the recommendations, but they were not apparent in the user interface that the buyers interacted with. And, over time the buyers came to expect that the recommendations included client segments. For example, towards the end of our study period, a buying director was listening to a presentation by the algorithms team and asked unprompted, "The client segment here is age?" The team of data scientists responded:

> **Data scientist 1:** Yes. For this group, it's four age segments. But that's changeable! For our longer-term vision, we can use the machine-learned personalized style for client segmentation.
>
> **Data scientist 2:** Yes. That's a very important vision for where we are headed.

Data scientist 1: We will use the best client segment that gives us the best outcomes

The team further elaborated their approach. One data scientist said, “We keep it away from users. We keep it ‘behind the scenes’ which means we could expand it to way more client segments” (i.e., we could add a lot of complexity without that additional creating complexity for the users). One of the merchandising team members asked, “So it could be multi-dimensional?” and the data scientist confirmed, “Yeah, whatever, however we want to segment that.” Putting these structuring decisions “behind the scenes” meant there was considerable flexibility in how the structuring happened. The buyers would still get recommendations that were informed by this extra layer, but they would not be aware of the complex structuring based on client segmentations that was happening “behind the scenes” to inform their recommendations. The data science team or the buying directors could select any client segmentation to use to inform the recommendations and could also choose to change those segmentations frequently.

As the buyers and data scientists envisioned further changes to the buyers’ roles beyond these changes of explicitly measuring the impact of their decisions and including more and including more complex segmentations “behind the scenes” in the recommendations that they received, they used language of buyers “curating” an algorithmically-recommended assortment, rather than producing the assortment plan themselves. They expected that the buyers would start to curate the algorithmic assortments for both context and for strategy. We saw examples of this curation role when we returned to observe the buyers use the tool. They ran the tool, projected the algorithmic recommendations on the screen and then as a group discussed removing different recommendations. For example, one set of styles were removed because one of the buyers knew that they were trying to reduce their work with a certain vendor. As another example, another set of styles was swapped out because one of the buyers thought that their groups’ vision for that season involved more of a pastel palette.

The buyers' role relationship with the planner role. Another change emerged in the buyers' role relationship with the planners. Recall that the data scientists always wanted to link the decision structuring to outcomes. They configured the assortment tool to allow the buyers to automatically model and compare different assortment plans. The buyers could input many ways of dividing out their buys (i.e., structuring their decisions) and easily "run" the optimization algorithm to get different sets of recommendations. In the past, working in the spreadsheets, the buyers tended to only have one assortment plan in mind. The planners and buyers shared ownership of the spreadsheets, and the planners' job was to model the buyers' emerging assortment plan and determine if they could make it work by moving the buys around. Within this role relationship, and using the more static spreadsheets, it would have been extremely complex to develop several different assortments plans and dynamically compare them.

One of the buying directors said in a cross-functional meeting, "Oh… this means we can now compare assortments?" and an algorithms manager confirmed: "Right, now we can compare assortments." The tool interface was also visually appealing, and let the buyers visualize their assortment as they were planning it. The buyers visualized the assortment on many different dimensions using these various product categories using a dynamic pivot table that displayed the different styles in small, medium, and large bubbles (see Figure 7). This new data and the tool interface allowed the buyers to sort and visualize their planned inventory many ways across many product dimensions. Thus, even though the tool was very deliberately designed to support one workflow within the buyer "leaf nodes," it ended up automating a lot of the manual computation that the adjacent planner role had previously done.

Insert Figure 7 here

When we returned to AlgoCo a year after our first study period, we observed a buying team go through an assortment planning process. They used the tool to easily and dynamically model and curate algorithmic recommendations. They were able to save different versions of the assortment plan and compare them. And the age segmentation was running in the background. This functionality of dynamically modeling the assortment plan, including client segments, was not possible in traditional spreadsheets, where the planners would have had to do extensive manual work to calculate the full set of targets for each possible assortment. There, any changing of units between different product categories – for example, between vendors or sizes – was manual and laborious and it tended to be done more in trading units between styles or different product categories. The new tool included as many product categories as were requested and the powerfully flexible UI allowed the buyers to dynamically model many ways of dividing up their units to model, compare, and choose among optimized assortment plans.

Reconfiguring the professionals' organizational hierarchy

AlgoCo also worked to reconfigure the buyers' organizational hierarchy as they came to see the issues that were created by the way the org chart constrained and influenced the algorithmic search space and related sets of recommendations. The buyers had structured the "people structure" hierarchy using practical, interpretable, and fairly static structuring – e.g., their people structure tended to be represented in typical PDF organizational charts that did not change very often. In contrast, the data scientists preferred flexible structuring that explicitly predicted better outcomes. They had come to understand that the org chart as visualized in the PDF also represented and constrained the way that the massive, centralized stores of data were stored and structured, as well as the way that the budget (including the assigned buys, targets, and metrics) was structured and allocated. Directors in both departments saw strategic opportunities to separating out these the different structures and more flexibly and dynamically modeling some of the decisions that were being constrained by the static org chart.

Work to decouple the data structure, budget structure, and org structure. Both departments took on many related initiatives trying to figure out how to support a more flexible and dynamic approach to structuring, especially regarding the intertwined data, budget, and people hierarchical structures. One of the data scientists led the work of conceptualizing this problem and of convincing other data scientists and AlgoCo leaders to work on decoupling the separate but related functions that were all structured around the same org chart decision tree. He gave a formal presentation focused on the data structure aspect where he explained,

> There's only one data structure hierarchy, and it's currently doing three things. Focus on two relevant things for now – “what is it” and “who bought it”. So, (*gesturing to the buying group level and related level in the data structure*) this hierarchical level is interpreted by us on algorithms as meaning something about “what is it” – “oh, it is a women’s blouse.”
>
> But what it really is really telling us, is actually “who bought it” – “oh this was bought by the blouses buying group.”

His description was explaining that the data structure was in fact encoding the people structure, rather than recording properties of the items themselves. To say it a different way, if anyone looked in the data, they would find that all the clothing item IDs were nested in product category IDs that mirrored the buying teams. But as described above, the segmenting out of the buying teams was arbitrary and based on manageability and interpretability, so there was no reason to use those hierarchical divisions to structure the data.

His proposal was to restructure the massive data store so that it was not hierarchically segmented at all. He said in one of his presentations where he was explaining the problem to other data science leaders, “I'm going to argue that not all of these things (the data, budget, and people structures) are hierarchical in nature, in fact, I think only one of them is." He argued that it was much more consistent with data science approaches to store the data as "flat" and use flexible "roll ups" instead of static hierarchical divisions. He explained:

> So, the proposal is to build a new data model where the principles are ... that we want to use hierarchy only for concepts that are truly hierarchical. When a hierarchy is not unambiguous, *tags are better than hierarchy*.
>
> My example here is, in old email clients there would be folders, and sometimes you could have folders within folders and that is like a hierarchical way of grouping your emails. If you had an email from your dad about buying a house, you would have to decide, "Does this go in the family folder or does this go in the real estate folder?"
>
> Then with Gmail, you just put tags on there. You don't have to make these choices about where does it go; you just tag it with everything that's relevant.

The email example helped explain his vision for a better way of structuring their data if they decoupled it from the way the merchandising org chart was structured. Every item would be simply tagged with as many relevant tags as desired. And then "roll ups" could aggregate all relevant items using tags depending on how the data were being used in that moment. In the email "tag" example, one could easily look at every email that had been tagged "family" and then separately look at every email that had been tagged "real estate"—the same email could easily appear in both. In contrast, when the email had been stored in a hierarchical folder structure, you could only see and understand the email as a family email or a real estate email. His point was that, in following the organizational chart, the data structure was only letting people see "women's blouse" rather than letting that item be tagged and flexibly rolled up into sets such as "any green item" or "anything from Vendor A" or "anything for millennial clients".

The plan to change the data structure to a flatter and more flexible model where all items were tagged rather than hierarchically stored was a huge undertaking, but also well-received within both data science and merchandising. One of the buying directors explained it this way in a meeting, "We are thinking about breaking the dependency of the data structure hierarchy and how Merchandising organizes themselves to allow for more flexibility…" The data scientists saw the flexibility in terms of the different analyses that could be done, and the buying directors saw opportunities in terms of how the merchandising teams were staffed and structured.

Roll up the leaf node. The data science directors and buying directors were focused on those strategic departmental structuring issues involving separating out the data structure from the merch department structure. The data science team who had specifically developed the assortment planning algorithm were also working on related issues connected to the merch organizational hierarchy. Recall that they saw the org chart as a decision tree and considered the structuring of the leaf nodes (the buyers' jobs where the decisions were made) as arbitrary and as unnecessarily constraining the search space in ways that were impacting outcomes. One phrase that caught on referred to the idea to "roll up a leaf node" and run the optimization recommendation algorithm there. Figure X visualizes what was meant by this idea. The original configuration of the algorithmic tool was to produce an optimized set of recommendations at one individual buyer's set of buys. The individual buyers were the "leaf node" of the decision tree and several buyers were together nested under a shared manager. "Rolling up the leaf node" meant aggregating all the buys and targets of an entire team of buyers and running the optimization algorithm across that level of buys (see Figure 8). This idea was the specific way of allowing the algorithm to "explore a larger space for better results." The data scientists proposed "rolling up the leaf node" on the plus buying team. This proposal meant that the manager and leaf node structure (typical org chart) would be reconfigured into a buying group that collectively curated the whole plus assortment. The algorithmic tool would model many Plus-wide assortments that could be compared, and the buying group would curate those group-level recommendations for context and strategy. The Plus buying director was willing to try this experiment and learn from the process of group-level algorithmic recommendation and group-level curation, which represents a significant shift from the traditional manager and leaf node way structure.

Insert Figure 8 here

As the data scientists worked on this idea of rolling up the leaf node and recommending and curating at that buying group level, they also started to think through and model other ways that the decisions could be structured. As an example, the data science team kept on their team roadmap charter the question of "planning at different levels of hierarchy" – which referred to all of the different ways they could learn from "rolling up the leaf node". They kept a brainstormed list of all the ways to do this, including "Department, Class, Silhouette, etc." One of the data scientist's strategic idea was to roll up decisions by client segments and organize the buyers into groups around the client segments. She explained, "It kind of makes sense to me to have buying groups organized around client segments" because client segments predicted variance in outcomes. These continuing discussions at AlgoCo about how they reconfigure the people structure, the data structure, and the data tools is well-summarized by some educational materials that one of the data scientists put together:

> When it's important to have the benefits of splitting finely while focusing on a small number of relevant segments, this is a great opportunity to let humans and machines do what they each do best.
>
> Algorithms can be designed to segment the data to as fine a granularity as the data supports. Then, algorithms can dynamically detect the key outlier segments that deserve extra attention, and aggregate the other segments back into larger groups.
>
> What gets surfaced to humans are the important findings about the forest, as well as highlights of the handful of trees that matter right now. Such flexible segmentation

> schemes enable people and algorithms to adapt together to changing data and changing business priorities.

The data scientists recognized the value of the more bureaucratic approaches that are practical and interpretable for people's decision-making, and also the algorithmic approaches to dividing and aggregating decisions and outcomes. Their aim was to flexibly balance these approaches going forward, which as we have shown involved and will continue to produce many changes to role structures and organizational hierarchies.

DISCUSSION

This 10-month ethnography of a retail technology company analyzes the work of retail assortment planning before and after the development of a new algorithmic tool. The comparison reveals why data scientists' work to develop algorithmic tools is producing significant changes to professionals' roles, role structures, and organizational hierarchies, and with what implications. Our data illustrate that the professionals' original structuring approach involved *bureaucratic decision structuring* to coordinate the collective decision-making of a 200-person department. This approach, which involved structuring a profession's organizational chart as a largely static "decision tree", was in tension with the *algorithmic decision structuring* enabled by the new algorithmic tool, which involved dynamically measuring the impact of decisions on outcomes. To resolve these tensions to effectively use of the algorithmic tool, the organization reconfigured roles, role structures, and hierarchies –significant structural organizational changes consistent with theories of major technological change (Barley, 1990; Barley, 2015; Barley, 2020).

Algorithmic disruption of bureaucratic decision structuring

Our society is on the brink of a myriad of different so-called technological "revolutions".

From the Big Data Revolution (e.g., Kitchin, 2014) to the Fourth Industrial Revolution (e.g., Schwab, 2017) to the Control Revolution (e.g., Barley, 2020), technology is changing society in fundamental ways. Yet new, transformative technology is not sufficient for spurring a revolution. In order for a technological revolution to truly be a revolution, it must involve a fundamental change in the technological infrastructure of a society (Barley, 2020). Algorithms constitute such an infrastructural change. As they are increasingly incorporated into organizational routines, algorithms are influencing authority (Pasquale, 2015), power dynamics (Kellogg, Valentine, & Christin, 2019) as well as decision making (Shrestha, Ben-Menahem, and von Krogh, 2019). It is critical for scholars and practitioners to understand how this algorithmic revolution is fundamentally changing the bureaucratic organizations that are at the core of modern society.

Nearly a half-century ago, Simon (1973: 272) asserted, "[t]o design effective decision-making organizations...we must understand the decision-making tools at our disposal, both human and mechanical". More than ever before, the decision-making tools at our disposal are in tension. It is critical to understand the nature of these tensions and how they might be resolved in order to design effective organizations. Whereas bureaucratic decision-making, as depicted on organizational charts, has long been considered rational (Weber, 1947), algorithms are "supercarriers of formal rationality" and offer "augmented rationality" (Cohen, 2007: 504; Lindenbaum, Vesa, and Den Hond, 2020: 250-259; Glaser, Valadao, and Hannigan, 2020). In this study, we show how the tension between the two forms of rationality—in the case of bureaucratic decision making, rationality based on human interpretability, and in the case of algorithmic decision making, 'advanced' rationality that is unconstrained by human interpretability—offers an as yet undertheorized reason why algorithms are bringing about large scale infrastructural change.

This study, thus, contributes to research on bureaucracy and organizational design. A key premise of bureaucratic decision-making is the belief that expertise is critical (Kweit and Kweit, 1980). Bureaucracies are structured to guide and "factor" decisions to effectively leverage local expertise. Jobs are nested in roles, which are "systems of prescribed decision premises" (Simon, 1991), that are constrained by rules and conceptions of controls. We show how, by facilitating top-down decision-making, bureaucratic decision structuring is fundamentally at odds with algorithmic decision structuring. Thus, contrary to the assumptions inherent in prior work (e.g., MacKenzie, 2014; Christin, 2018; Brayne and Christin, 2020), we illustrate how bureaucratic structuring may hinder the effective structuration of expertise in bureaucratic organizations.

Contributions to theories of the structuring of professional expertise

This study also contributes to research on professions, addressing recent calls for a reevaluation of the professions literature to account for the changing nature of work (e.g., Anteby et al., 2016). Although scholars initially characterized "the profession" and "the bureaucracy" as "antithetical" (Freidson, 1984: 10), they have since acknowledged that they are, in fact, compatible. This compatibility involves professionals' roles being nested in managerial hierarchies through vertical integration and bound by the controls that distribute decision-making responsibility. As professional expertise is being increasingly structured in bureaucracies (Huising, 2015; Anteby, Chan, and DiBenigno, 2016), we have reached a critical moment in time wherein organizations are moving towards automating bureaucratic structures and controls. We show why and how the bureaucratic structuring of professional expertise is likely to drive structural changes as algorithmic decision structures are increasingly implemented in organizations.

To date, research has assumed that the expansion of professional expertise over particular tasks and decisions occurs through jurisdictional claims, often against other occupations, that are typically competitive and adversarial in nature. Our research addresses a yet unanswered key question related to the professions: "How do occupational members actually coordinate with a network of relations to collectively expand their scope of expertise?" (Anteby, Chan, and DiBenigno, 2016: 220). We show how algorithms can redefine role relations and encourage cooperative interaction between professional groups that have traditionally been separated in bureaucratic organizations by jurisdictional claims. By occasioning new divisions of labor, role relations, and ways of decision making, we show how algorithms can change the ways in which professions are structured in organizations and expand professionals' scope of expertise, such as buyers and planners newly working together to expand their expertise in creating a high-performing inventory that accounts for non-practical and non-human-interpretable categories.

Third, we contribute to research on algorithms and professions. While organizations still bear the imprints of prior conceptions of control, especially financial controls and related metrics, they are being increasingly dominated by quantification and algorithms as metrics of control (Fligstein, 1991; Espaland and Stevens, 2008; Kellogg, Valentine, and Christin, 2020). A rich recent body of work has shown that quantification and algorithms change the nature of professional expertise and, in particular, professionals' jurisdictions and areas of discretion (e.g., Christin, 2017). As has been the case throughout the rise of bureaucratic structures, we show how algorithms as new conceptions of control are not only changing the nature of professional expertise, but also how professionals are structured in organizations.

We also illuminate why, in addition to understanding the "black box" or opaqueness that is characteristic of algorithms (Faraj, Pachidi, and Sayegh, 2008), it is also critical for scholars and practitioners alike to understand the "black box" of bureaucratic decision structuring—the often taken-for-granted structures that are fundamental to bureaucratic structures. By gaining a heightened understanding of how bureaucratic organizations structure decisions, practitioners can better appreciate how to increase the adoption of algorithmic tools to resolve key tensions, while scholars can adopt a new lens through which to study the implementation of algorithms. While algorithmic technologies may be occasions for structuring (Barley, xx), that structuring also depends on the bureaucratic structures that they may be aiming to displace.

Contributions to theories of organizational design

We also contribute to research on organizational design within the Carnegie School tradition. Within this classic body of work, the information processing point of view holds that effective organization design involves a division of labor that is able to factor the total system of decisions that need to be made into relatively independent subsystems that "can be designed with only minimal concern for its interactions with the others" (Simon, 1976). As Simon (1957) explained, the division of labor in organizations has been necessary because humans' processing power and attention are limited. We show how algorithms may fundamentally change what we consider to be ideal organizational forms as algorithmic technology enhances humans' processing power and supports 'arbitrarily many' decisions.

Building on Simon's (1962) assertion that complex organizations perform better when they have a hierarchy and "near decomposable structures," a rich body of work has investigated modular organizational designs, often in the context of M-form organizations (e.g., Galunic and Eisenhardt, 2001). Yet, with modular systems, there is a tradeoff between the breadth of search and the speed of search (Ethiraj and Levinthal, 2004). While prior research has shown that organizations leverage modular forms to realign organizational structures to an evolving mix of product markets by competing for resources and charters, we show how organizations leveraging algorithmic decision structuring may fundamentally change role structures to enable the collaborative pursuit of new charters.

This study offers a contrast to prior work that has shown how algorithms enable new ways of coordination by managing task decomposition, sometimes as part of a divide-and-assign strategy (Zammuto et al., 2007; Faraj, Jarvenpaa, & Majchrzak, 2011; Boudreau & Lakhani, 2013). Our research shows how, while “organizing no longer needs to take place around hierarchy” (Zammuto et al., 2007: 749), it, in large number, still does, with the implication that the impact of bureaucratic structuring on algorithmic structuring ought to be considered.

“Moneyball” for the organizational chart

As we conclude the current paper, we note a connection to a well-known example of data analytics changing the way that decisions were made, as popularized in the book-turned-movie, Moneyball, by Michael Lewis. We came to refer to this paper as “Moneyball for the org chart” and found that metaphor informative, so wanted to draw that connection here. Lewis (2004) traces the introduction of “sabermetrics” to baseball, which involved the use of statistical methods to evaluate player performance. Prior to the use of sabermetrics—originally defined as “the search for objective knowledge about baseball” (Hirsch & Hirsch, 2014), baseball scouting was based on methods of intuitive categorization of players. Scouts were instructed to seek out players based on several categories of capabilities—for example, strength or a short stride—that intuitively reflected the characteristics of a great baseball player, especially his power. Scouts used a 2-to-8 scale to quantify the value of a player, with 5 representing the average skill of a major-league player. With the use of sabermetrics, the methods of categorization and commensuration changed. Lewis recounts how Billy Beane, the General Manager of the Oakland A’s, started to develop new metrics such as on-base percentage (OBP) and the slugging percentage that did not emphasize the physical power of an athlete and did not correspond to human interpretable categories. As well, with Moneyball, staffers began to structure roles differently. They no longer tried to pick the best individual players, but instead, tried to seek the best aggregation of players for the money by determining which players contributed more to winning than their salaries suggested they could. Thus, they started to select for wins, rather than individual player performance. An analogy can be drawn to the current study and our findings that show how new analytical techniques fundamentally transformed how role relations and hierarchies were structured at AlgoCo, enabling its members to dynamically make decisions based on outcomes, rather than human-interpretable roles as structured by the organizational "roster", otherwise known as the org chart.

References

Abbott, A.
1988 The System of Professions: An Essay on the Division of Expert Labor. Chicago: University of Chicago Press.

Abernethy, M. A., and J. U. Stoelwinder
1995 "The role of professional control in the management of complex organizations." Accounting, Organizations and Society, 20: 1-17.

Anteby, M., C. K. Chan, and J. DiBenigno
2016 "Three lenses on occupations and professions in organizations: Becoming, doing, and relating." The Academy of Management Annals, 10: 183-244.

Bailey, D. E., and S. R. Barley
2020 "Beyond design and use: How scholars should study intelligent technologies." Information and Organization, 30: 100286.

Barley, S. R.
1986 "Technology as an occasion for structuring: Evidence from observations of CT scanners and the social order of radiology departments." Administrative Science Quarterly, 31: 78-108.

Barley, S. R.
1990 "The alignment of technology and structure through roles and networks." Administrative Science Quarterly, 35: 61-103.

Barley, S. R., and P. S. Tolbert
1991 "Introduction: At the intersection of organizations and occupations." Research in the sociology of organizations, 8: 1-13.

Barley, S. R.
2015 "Why the Internet makes buying a car less loathsome: How technologies change role relations." Academy of Management Discoveries, 1: 5-35.

Barley, S. R.
2020 Work and Technological Change. Oxford: Oxford University Press.

Berlant, J. L.
1975 Profession and Monopoly: A Study of Medicine in the United States and Great Britain. Berkeley, CA: University of California Press.

Bourgeault, I. L., K. Hitschkorn, and I. Sainsaulieu.
2011 "Relations between professions and organizations: More fully considering the role of the client." Professions and Professionalism, 1: 67-86.

Boudreau, K. J. and K. R. Lakhani

2013 “Using the crowd as an innovation partner.” Harvard Business Review, 91: 61–69.

Briscoe, F.
2007 "From iron cage to iron shield? How bureaucracy enables temporal flexibility for professional service workers." Organization Science, 18: 165-335.

Brayne, S., and A. Christin
2020 “Technologies of crime prediction: The reception of algorithms in policing and criminal courts.” Social Problems, spaa004.

Bunderson, J.S., S.M. Lofstrom, and A.H. Van de Ven.
2000 “Conceptualizing and measuring professional and administrative models of organizing.” Organizational Research Methods, 3: 366–91.

Christin, A.
2018 “Counting clicks: Quantification and variation in web journalism in the United States and France.” American Journal of Sociology, 123: 5.

Christin, A.
2017 "Algorithms in practice: Comparing web journalism and criminal justice." Big Data & Society, 4.

Cogan, M. L.
1955 “The problem of defining a profession.” The ANNALS of the American Academy of Political and Social Science, 297: 105–11.

Cohen, M. D.
2007 “Perspective—Administrative Behavior: Laying the foundations for Cyert and March.” Organization Science, 18: 503-506.

Cooley, M.
1999 “Human-centered design.” In R. Jacobson (eds.), Information Design: 59–83. Cambridge, MA: MIT Press.

Cyert, R. M., and J. G. March.
1963 A Behavioral Theory of the Firm. Englewood Cliffs, NJ: Prentice-Hall.

Edmondson, A. C., and S. E. McManus.
2007 “Methodological fit in management field research,” Academy of Management Review, 32: 1155–1179.

Espeland, W. N., and M. L. Stevens.
2008 "A sociology of quantification." European Journal of Sociology, 49: 401-36.

Ethiraj, S. K., and D. Levinthal.
2004 “Modularity and innovation in complex systems.” Management Science, 50: 159–173.

Faraj, S., S. Pachidi, and K. Sayegh.
2008 “Working and organizing in the age of the learning algorithm.” Information and Organization, 28: 62-70.

Faraj, S., S. L. Jarvenpaa, and A. Majchrzak.
2011 “Knowledge collaboration in online communities.” Organization Science, 22: 1224-1239.

Freidson, E.
1984 “The changing nature of professional control.” Annual Reviews, 10: 1-20.

Fligstein, N.
1990 The Transformation of Corporate Control. Cambridge, MA: Harvard University Press.

Freidson, E.
2001 Professionalism, the Third Logic: On the Practice of Knowledge. Chicago: University of Chicago Press.

Galunic, D. C., and K. M. Eisenhardt.
2001 “Architectural innovation and modular corporate forms.” Academy of Management Journal, 44: 1229-1249.

Glaser, V. L., R. Valadao, and T. R. Hannigan.
2020 “Algorithms and routine dynamics.” Cambridge Handbook of Routine Dynamics.

Goode, W. J.
1957 "Community within a community: The professions." American Sociological Review, 22: 194-200.

Hall, R. H.
1968 "Professionalization and bureaucratization." American Sociological Review, 33: 92-104.

Hinings, B.
2006 ‘‘The changing nature of professional organizations.’’ In S. Ackroyd, R. Batt, P. Thompson, and P. S. Tolbert (eds.), The Oxford Handbook of Work and Organization: 404–424. Oxford: Oxford University Press.

Hirsch, S., and A. Hirsch.
2014 The Beauty of Short Hops: How Chance and Circumstance Confound the Moneyball Approach To Baseball. Jefferson, North Carolina: McFarland & Company.

Huising, R.
2015 "To hive or to hold? Producing professional authority through scut work." Administrative Science Quarterly, 60: 263-299.

Kellogg, K. C., M. A. Valentine, and A. Christin.

2019 “Algorithms at work: The new contested terrain of control.” Academy of Management Annals, 14: 366–410.

Kitchin, R.
2014 The Data Revolution: Big Data, Open Data, Data Infrastructures and Their Consequences. Thousand Oaks, CA: Sage Publications

Kweit, R. W., and M. G. Kweit.
1980 "Bureaucratic decision-making: Impediments to citizen participation." Polity 12: 647-66.

Larson, M. S.
1977 The Rise of Professionalism: A Sociological Analysis. Oakland, California: University of California Press.

Lebovitz, S., Hila L.A., and N.Levina
2020 “To incorporate or not to incorporate AI for critical judgments: The importance of ambiguity in professionals’ judgment process.” Working Paper.

Lewis, M.
2004 Moneyball: The Art of Winning an Unfair Game. Manhattan: W. W. Norton & Company.

Lindenbaum., D., M. Vesa, and F. D. Hond,
2020 “Insights from “The Machine Stops” to better understand rational assumptions in algorithmic decision making and its implications for organizations.” Academy of Management Review, 45: 250-259

MacKenzie, D.
2014 “A sociology of algorithms: High-frequency trading and the shaping of markets.” University of Edinburgh Working Paper.

MacKenzie, D.
2018 "Material signals: A historical sociology of high-frequency trading." American Journal of Sociology, 123: 1635-1683.

Norman, D. A.
2005 “Human-centered design considered harmful.” Interactions, 12: 14–19.

Noordegraaf, M.
2011 "Risky business: How professionals and professional fields (must) deal with organizational issues." Organization Studies, 32: 1349-1371.

Parry, N., and J. Parry
1976 The Rise of the Medical Profession: A Study of Collective Social Mobility. Oxford: Routledge.

Sachs, S. E.

2019 “The algorithm at work? Explanation and repair in the enactment of similarity in art data, Information, Communication & Society,” 1689-1705.

Schwab, K.
2017 The Fourth Industrial Revolution. London: Penguin Books Limited

Shrestha, Y. R., S. M. Ben-Menahem, and G. V. Krogh,
2019 “Organizational decision-making structures in the age of artificial intelligence.” California Management Review, 61: 66-83.

Simon, H.A.
1957 Models of man; social and rational. Hoboken: Wiley.

Simon, H. A.
1962 “The architecture of complexity.” Proceedings of the American Philosophical Society, 106: 467–482.

Simon, H. A.
1973 "Applying information technology to organization design." Public Administration Review 33: 268-78.

Simon, H. A.
1976 Administrative Behavior. New York City: Macmillan Inc.

Simon, H. A.
1991 “Bounded rationality and organizational learning,” Organization Science, 2: 125–134.

Parsons, T.
1968 "Professions." International encyclopedia of the social sciences, 12: 536-547.

Pasquale, F.
2015 The Black Box Society: The Secret Algorithms That Control Money and Information. Cambridge, MA: Harvard University Press.

Van Ossel, G.
2019 Optichannel Retail. Beyond the Digital Hysteria. Develop and Implement a Winning Strategy as a Retailer or Brand Manufacturer. Belgium: Lannoo Publishers

Weber, M.
1947 The Theory of Social and Economic Organization. New York: Free Press.

Weber, M.
1958 The Protestant Ethic and the "Spirit" Of Capitalism and Other Writings. London: Penguin Classics.

Wilensky, H. L.

1964 “The professionalization of everyone?” American Journal of Sociology, 70: 137-158.

Zammuto, R. F., T. L. Griffith, A. Majchrzak, D. J. Dougherty, and S. Faraj
2007 ‘‘Information technology and the changing fabric of organization.’’ Organization Science, 18: 749–762.

Table 1. Comparison of Bureaucratic and Algorithmic Decisions Structures

	Bureaucratic Decision Structures	**Algorithmic Decision Structures**
Number of decision structures in use	Typically one at a time	"Arbitrarily many"
Criteria for deciding between decision structures	Interpretability Feasibility	Predicts variance in data Optimized outcomes
Relevant technologies	● PowerPoints and PDFs with images of organizational charts ● Excel spreadsheets with Macros	● Dynamic user interface that allows for modeling and curation ● Backend that integrates many data and algorithmic systems (e.g., Figure 1)
Models of change	Engage in traditional change management: ● update all PowerPoints and spreadsheets ● change reporting relationships and groupings of individuals	Local users learn to use dynamic decision structures: ● users learn to run different scenarios, ● data scientists learn to display different information Organization learns to accommodate disrupted bureaucratic decision structure (Barley, 2020) ● reimagine roles ● reimagine hierarchy ● reimagine accountability

Table 2. Reconfigured Roles, Role Structures, and Professional Hierarchies

	Original **Using bureaucratic decision structuring**	**Reconfigured** **To accommodate algorithmic decision structuring**
Buyers' Role	Buyers use non-standard and practical product categories to guide decision making. Buyers' intuition and impact of decisions on outcomes is not explicitly measured. Buyers' role involves being an artist with intuition for upcoming trends.	Buyers visualize impact of non-standard and practical product categories using tool that accommodates various constraints. Buyers' intuition and impact of their decisions on outcomes become recorded, measured, and analyzed over time. Buyers' role involves becoming a curator of algorithmically-recommended assortment, curating for context and strategy.
Buyer-Planner Role Structure	Buyers 'hand off' vision for well-performing inventory to planners. Planners roles' involves carrying out buyers' visions and making numbers work. Clear division of labor between buyers and planners.	Buyers begin to dynamically compare multiple assortment plans and, in turn, make decisions that previously involved planners' manual computation. More flexible division of labor between buyers and planners.
Buyer Professional Hierarchy	Hierarchy is structured using practical, interpretable, and static categories. Buyers' jobs are "leaf nodes" where decisions are made, with clear and codified managerial hierarchy.	Manager and leaf node structure reconfigured into a buying group that collectively curates entire assortment. New flexible hierarchies are considered that enable buyers and algorithms to adapt together.
Buyer-Planner Hierarchy	Planning director sets top-down allocation of buys and targets and "sends them down" the static structured org chart.	Planning directors look to redefine their expertise in setting department wide strategies. New opportunities are identified in terms of how buyer-planner teams are staffed and structured.

Figure 1. Mockup of Data Scientists' Explainer of their 'Algo' in the Larger System

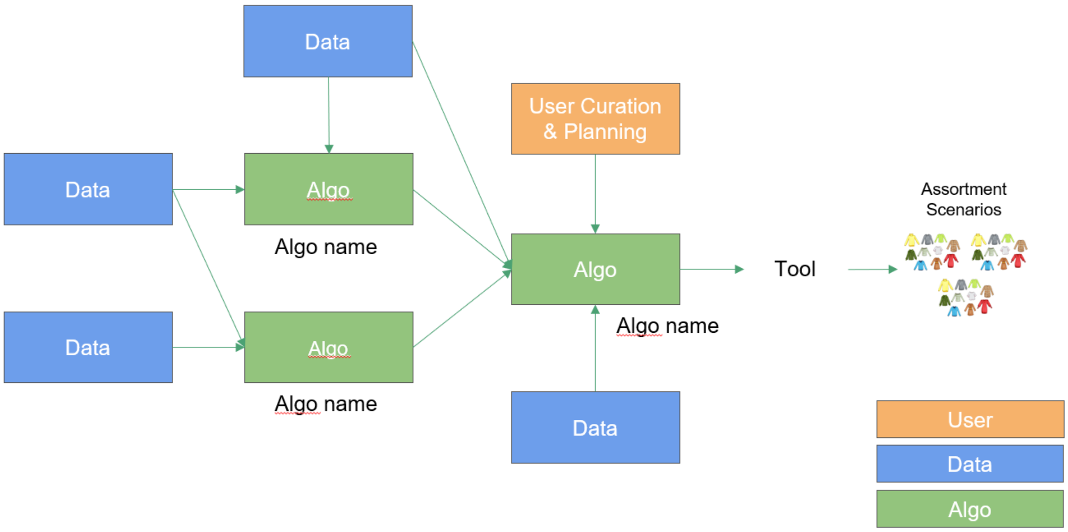

Figure 2. Mockup of Assortment Plan: Styles to Develop and Stock to Sell to Clients

Figure 2. Professionals' Organizational Chart (stylized) with non-standard categories

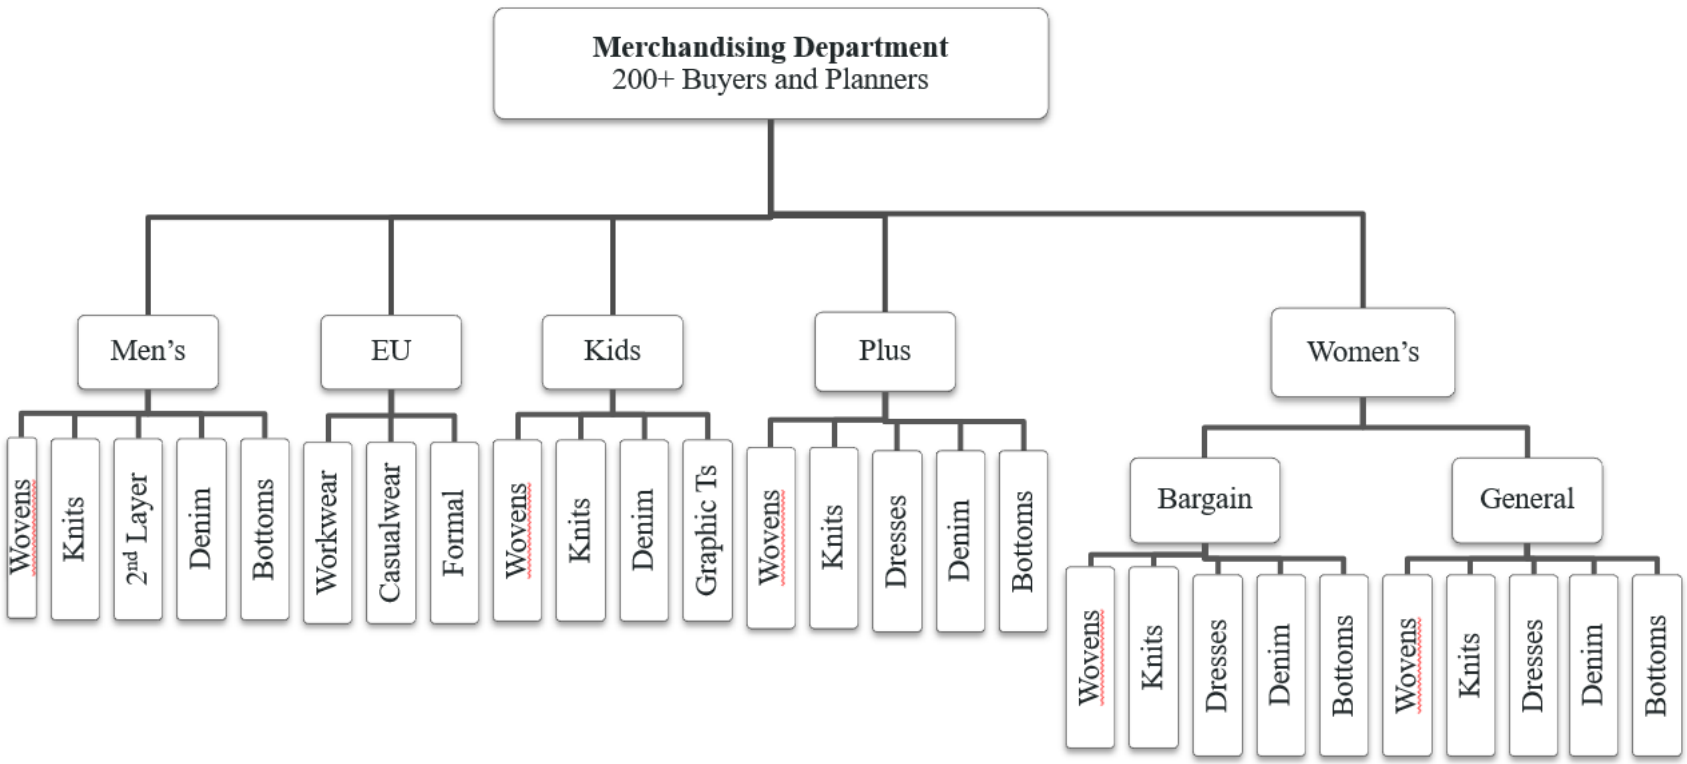

Figure 3. How Targets were Structured by Org Chart: Department to Buying Groups

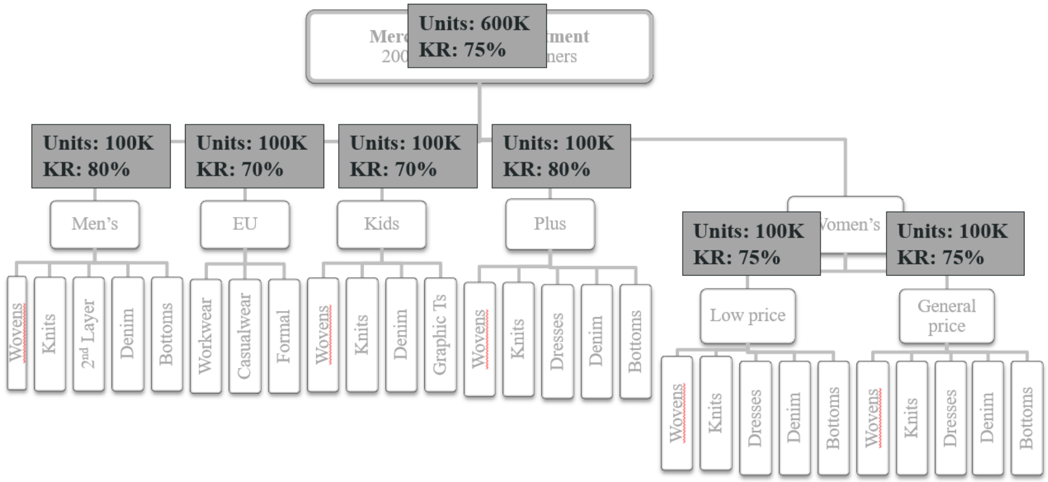

Figure 4 (con't). Buying Groups to Buyers

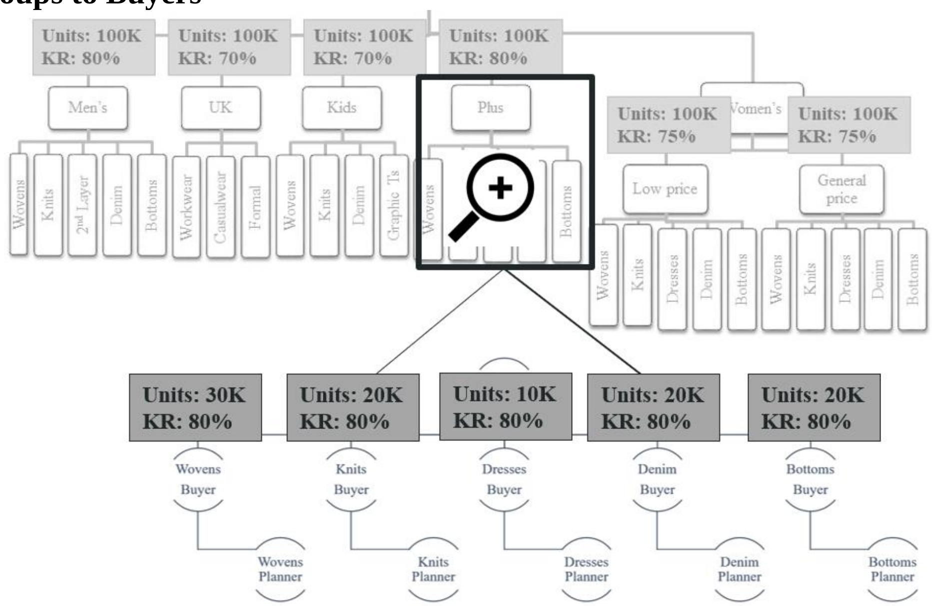

Figure 5. Example of Spreadsheets for a Buyer-Planner Pair

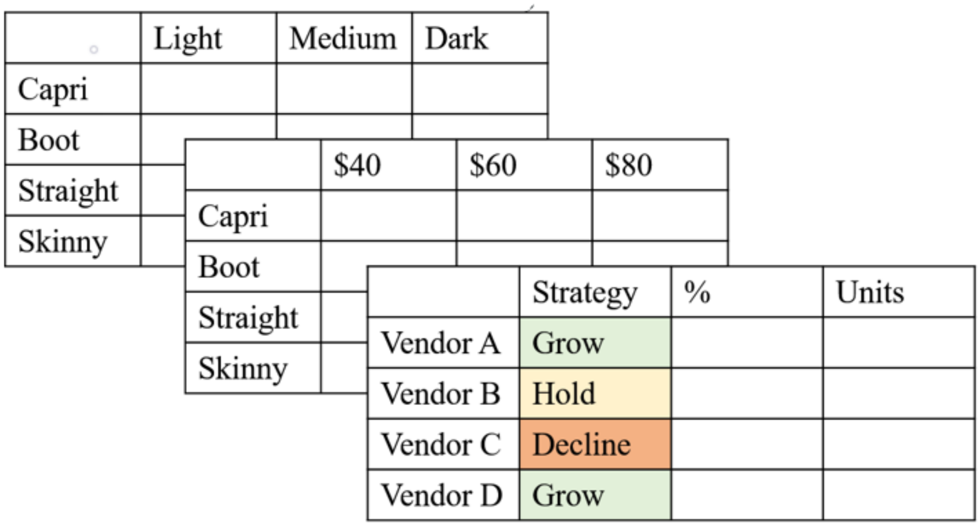

Buyers and planners would be given a set of buys that they would distribute across these spreadsheet cells to produce a diverse and well-performing inventory.

Figure 6. Including Client Segments for a Buyer-Planner Pair

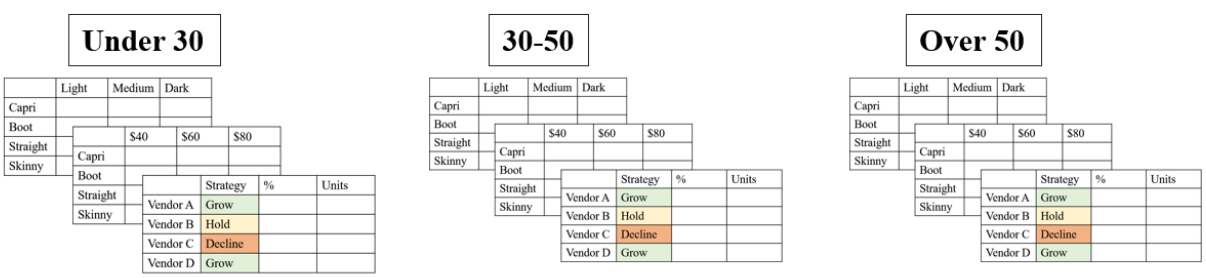

Including insights on client purchasing patterns in this process, when using the spreadsheets, meant multiplying the amount of work to calculate the predicted performance of the assortment

Figure 7. Mock-up of Algorithmic Tool User Interface

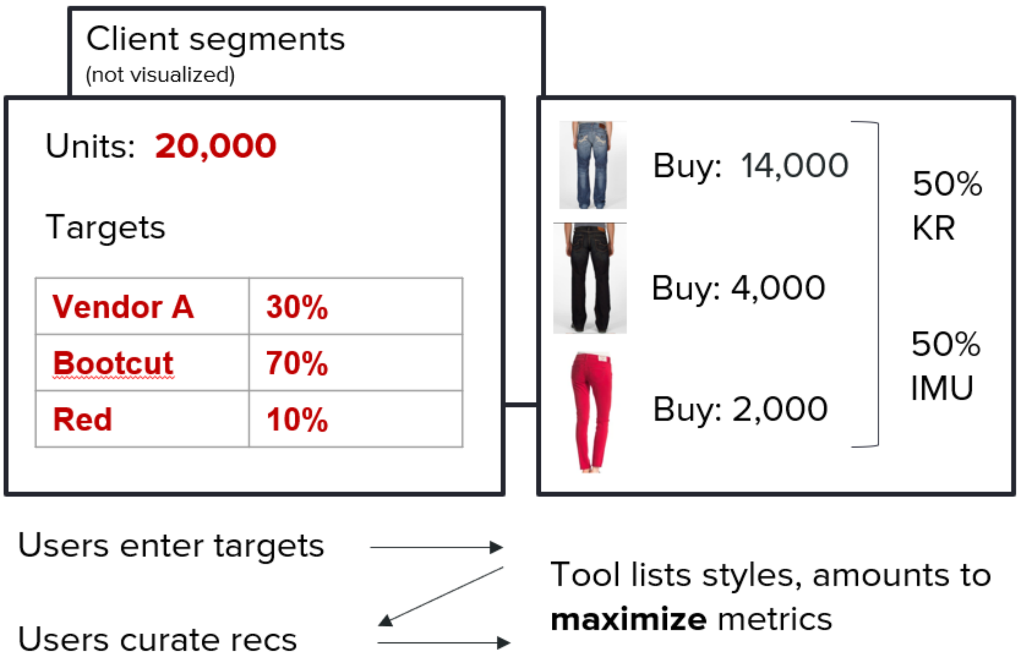

Figure 8. "Roll up the leaf node"