The Algorithm and the Org Chart: How Algorithms Can Conflict with Organizational Structures

ANONYMOUS AUTHOR(S)
SUBMISSION ID: 1517

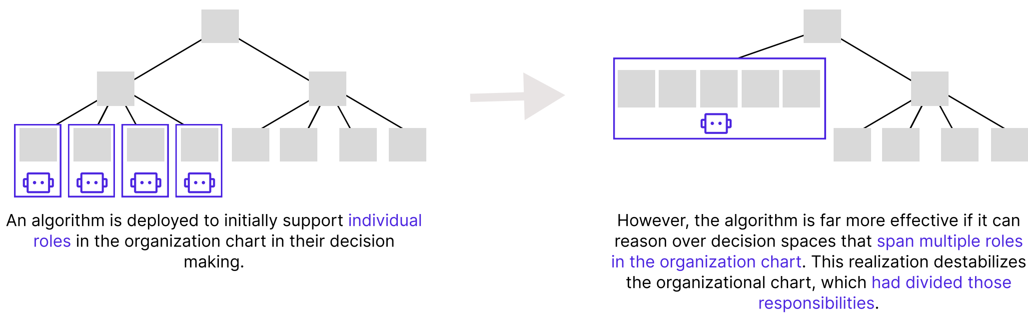

Fig. 1. A theory of the algorithm and the organization chart: while algorithms are often thought of as impacting individual workers' jobs, they can also come into tension with the core decision-making structures of an organization. This change arises when the decision spaces that produce the best algorithmic recommendations are in tension with the human decision spaces articulated in the existing organization chart.

Algorithms are introducing changes to individuals' jobs, but do algorithms also lead to changes in the structures of organizations themselves? Organizational structures, as often formalized into org charts, are meant to facilitate coordinated decision-making. Yet our 10-month ethnographic study of a large online retail company reveals why the organizational structures that facilitate effective decision-making by humans may be in tension with the organizational structures that facilitate effective decision-making using algorithms. Our findings show that the human decision-makers needed small, divided-up sets of decisions, and they had previously accomplished this in how they structured individuals' roles and teams in the org chart. In contrast, when data scientists developed a new algorithm and first deployed it within organizational structures meant to support human decision-making, they realized that these small divided-up decision spaces were arbitrarily constraining the algorithm's search space. When not constrained in this manner, the algorithm could identify and recommend better solutions, but those optimal solutions did not always align with the structure of roles and teams in the org chart. This study suggests that as algorithms are integrated into the workplace, organization designs may begin to more explicitly reflect the contours of those algorithms' behaviors.

CCS Concepts: • **Human-centered computing → Computer supported cooperative work**; **Empirical studies in collaborative and social computing**.

Additional Key Words and Phrases: algorithms, automation, planning, hierarchy, organizational structure, ethnography

---

Conference acronym 'XX, November 9–13, 2024, San José, Costa Rica

ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00
https://doi.org/XXXXXXX.XXXXXXX

ACM Reference Format:
Anonymous Author(s). 2023. The Algorithm and the Org Chart: How Algorithms Can Conflict with Organizational Structures. In . ACM, New York, NY, USA, 31 pages. https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

Algorithms, defined as "encoded procedures for transforming input data into a desired output, based on specified calculations" [29], are introducing profound changes to individuals' jobs, including their expertise, skills, and task boundaries, and to their managers' attempts at organizational control [e.g., 18, 26, 37, 38, 80]. The introduction of previous technologies, such as novel medical imaging modalities [5] and the global, digital communications systems [59, 72], affected individuals' roles and, in turn, changed organizational structures [3].

Given this history, could algorithms also lead to broader changes in the structures of the organization? And if so, why?

Most prior research focuses on the effects of algorithms on individual workers' decision-making: algorithms are configured to take on decisions and judgments typically considered knowledge work, or skilled work that is based on experts' critical thinking and decision-making [23]. This prior research focuses on experts making decisions within the purview of their individual jobs and has explored how and why they respond to new algorithms, given that algorithms may draw on different inputs, use different analytical processes, and sometimes produce different results and recommendations than the experts. As examples, prior studies have characterized changes to the individual decision-making of police officers, journalists, HR recruiters, radiologists, investment bankers, and retail fashion buyers [e.g., 2, 19, 22, 39, 71, 77, 78].

Yet key technology change theories predict that "transformative" technologies change more than just individual work practices within jobs — these technologies also impact broader organizational structures, including the ways that expert roles interact (i.e., their role structures) and their organizational hierarchies [3, 6, 7], both of which are central topics of interest to CSCW [e.g., 4, 8, 36, 64]. Roles, role structures, and organizational hierarchies are often codified in the organization chart, known colloquially as the "org chart." Changes to individual decision-makers' roles may easily ripple out to changing the organization chart as well. However, to date, there exists a gap in our understanding of how algorithms result in changes to broader organizational structures and, in turn, to the collaborative decision-making those structures are intended to coordinate.

In this paper, we report findings from a 10-month ethnographic study of a large online retail company where we encountered this issue while studying the development of a inventory planning new algorithm. During our study, data scientists developed a novel algorithm for the company's fashion buyers. Fashion buyers planned and purchased the large volumes of clothing inventory that the company stocked to sell to customers. The new algorithm was developed to recommend inventory plans to buyers based on historical sales data; it recommended a set of styles that would optimize buyers' assigned performance metrics. The algorithm was configured for the individual buyers' use but quickly came into tension with the entire organization chart of the Merchandising Department.

Previously, the organization chart had coordinated the work of planning the entire inventory by dividing up decisions among product segments that aligned with teams (e.g., the Plus Women's buying team, the Men's buying team) and sub-segments that aligned with roles on each team (e.g., the Plus Women's Denim buyer, the Men's Denim buyer). This style of organization chart has been used in retail companies for over a hundred years and optimizes for clear management and decision-making structures. However, as the buyers began to use the algorithm to recommend inventory plans for their sub-segments, they realized several issues with dividing up the inventory decisions into segments and sub-segments. Specifically, they saw that the algorithm would recommend different inventory plans depending on the decision space the algorithm was given (i.e., it would recommend a different plan if configured to recommend for the whole buying team vs. each individual buyer), and the algorithm could also measure which plan was more effective on a given set of performance metrics, relative to existing inventory plan metrics. These better-performing inventory plans were possible because the algorithm could identify many styles that were being missed when the decisions were segmented (e.g., masculine-styled clothes for women-identified customers), could recommend complementary styles that had previously been isolated within different roles, and could recommend more flexible changes in style distributions. However, the organization chart had disallowed such inventory plans because it segmented those decisions across roles - a segmentation that the organization began to reflect on and evolve in response to the algorithm's results.

Our findings analyze these tensions to contribute to a theory of algorithms and the organization chart: we show that the decision spaces which produce the best algorithmic recommendations may be in tension with the human decision spaces outlined in organization charts. This tension occurs because their objectives differ: an organization chart is designed to support human decision-making given humans' limited information processing and coordination capacities, whereas an algorithm is designed to maximize its objective regardless of those information processing and coordination needs - for better or for worse. Though we study the Merchandising Department of a retail company, this finding may extend to other companies where organization charts repeatedly divide targets and subsequent decision-making into human-sized decision spaces, including segmenting engineering departments around product lines, sales departments around geographies, and client services around industry targets. We suggest that there are two possible high-level outcomes from these tensions. One, detailed in prior work [11, 39, 40], involves workers undercutting, delegitimizing, or otherwise minimizing the impact of the algorithm in order to maintain the existing order. The other approach, as suggested by our ethnography, involves changing the organization chart to accommodate the algorithm's decision space: an integration of both algorithmic and human information processing.

2 RELATED WORK

In this section, we motivate why algorithms may come into tension with organizational structures through an integration of the literature describing algorithms' impact on individual decision-making and the organization theory literature relating coordinated decision-making to organizational structures.

2.1 Algorithms: Tools for Individual Decision-Making

Most studies of algorithms in the workplace have focused on the impact of algorithms on an individual's decision-making processes, or, at most, the processes of small teams. The primary thrust of much of this literature is to understand how humans and algorithms can best work together. Topics cover how advice generated by algorithms affects decision-making [e.g., 17, 27, 30], how the presence of algorithmic-agents affects perceptions of team attributes [e.g., 21, 49, 62, 63], what and how information provided to human decision-makers changes algorithm-supported decisions [e.g., 15, 34], and techniques for enabling algorithm-supported decision-makers to overcome barriers to superior human-algorithm performance [e.g., 14, 35, 48], such as aversion, overreliance [e.g., 13, 16, 75], and anchoring [e.g., 58].

Much of this work intersects with decision-makers' mechanisms for simplifying decision-making to accommodate limited processing capacities [65]. For example, Cai et al. [15] explore how medical experts' existing mental models created specific information needs within individual decision-makers when being trained to use clinical decision-support tools. With respect to overcoming decision-makers' overreliance, aversion, and anchoring, theorists often focus on the cognitive shortcuts which may underlie these behaviors. Rastogi et al. hypothesize that decision-makers anchor on the output of algorithms as a rational choice between time and accuracy and show that, in fact, increased cognitive resources in the form of more decision-making time does result in less anchoring [58]. Buçinca et al. [13] hypothesized that decision-makers over-rely on algorithms because decision-makers form heuristics about an algorithm's performance overall, rather than engaging with each prediction from the algorithm. Buçinca et al. go on to show that forcing functions, meant to disrupt heuristics-based thinking, do reduce overreliance [13]. Vasconcelos et al. [75] also show how human overreliance on algorithms is a result of rational decision-making under conditions of limited cognitive capacity or satisficing and that humans are less likely to over-rely on AI when the "cost," of checking the AI's output is low and the reward is high.

Organizations and their structures are also mechanisms that support decision-making by reducing information processing requirements and thus, the "cost" of making decisions [67]. And as such, organizations are likely to also be impacted by changes in algorithm use, given changes in decision-making at the individual level[e.g., 76]. But to date, organizations enter the literature on algorithms largely as mechanisms that constrain the development or use of algorithmic decision-making [e.g., 25, 52, 57]. For example, Rakova et al. [57] found that organizational contexts constrained the success of responsible AI initiatives.

However, little research directly studies how algorithms are impacting role structures and organizational hierarchies more generally, even though theories of technology change predict such effects are likely to unfold [7]. Surprisingly, this lack of research is not limited to the CSCW literatures, but extends into the domains of Organization Theory and Management of Information Systems. The lack of research may reflect the moment in time: typically individual work practices change before changes in role structures and organizational networks emerge [7, Figure 2.1]. Individuals' practices are often slow to evolve, relative to formal mandates to change, and changes to role structures and organizational hierarchies may occur even more slowly, after the close of many ethnographies, or may only be emerging within organizations more recently. The goal of this paper is to initiate development of this needed understanding. We draw on organizational theory as a useful explanatory literature to explore how algorithms may be informing such changes because of its focus on the structure of organizations and use of organization charts. We start with a review of how organization charts are used to structure coordinated decision-making among large groups of employees.

2.2 The Organization Chart: Organizations and Coordinated Decision-Making

According to organizational theorists, organizations are infrastructures for helping large groups process information and coordinate their decision-making. Organization charts, as the visualization of this infrastructure, facilitate this function by dividing up large sets of decisions into human manageable and interpretable domains (i.e., roles or jobs) and nesting them into organizational hierarchies that help coordinate across those domains [e.g., 4, 36, 64, 67]. Organizational structures facilitate coordinated decision-making among the large number of employees who comprise an organization in at least three ways.

First, organizational structures provide the blueprint by which decisions get divided into human-manageable quantities. These divisions are necessary because of the "bounded rationality of both humans and computers" [67, pgs. 240-241]. From an information processing perspective, an organizational hierarchy can be conceptualized as a series of "boxes-within-boxes" [68, p. 128] which factorizes decisions into sub-decisions. The most granular boxes in this decision hierarchy contain a number and size of decisions "reduced to manageable proportions" [67, p. 241]. The decisions in these boxes - along with related tasks - define a job or role within the organization that is achievable by a single human [9, 31, 43, 67]. According to functionalist theories of organizational design, the way that these roles are grouped within the organization hierarchy - both horizontally in teams and vertically in layers - informs and is informed by the level of coordination necessary between roles and signals appropriate lines of accountability [28, 31, 43, 66, 67, 70, 79].

Second, organizational structures provide contextual information that guides human decision-making and reduces the number of alternatives a decision-maker considers. Specifically, the social context of any role defines a "decision premise" that guides appropriate actions of the role [65, p. 201]. "Roles tell organization members how to reason about the problems and decisions that face them: where to look for appropriate and legitimate informational premises and goal (evaluative) premises, and what techniques to use in processing these premises" [66, pgs. 126-127]. In this way, organizational structures provide a certain environmental context for individuals, reducing the alternatives individuals will consider in their decision-making, and decreasing the information processing necessary when enacting a role.

Finally, organizational structures outline repeated patterns of activity between and among group members. According to Galbraith, "Planning achieves integrated action and also eliminates the need for continuous communication among interdependent subunits as long as task performance stays within the planned task specifications, budget limits and within targeted completion dates...the ability of an organization to coordinate interdependent tasks depends on its ability to compute meaningful subgoals to guide subunit action." [28, p. 29]. Mintzberg argues that the definitions, decision-premises, delineations, organizational position, and coordinating mechanisms of and between roles change relatively infrequently [45, p. 86] and become "givens" [43, 67]. These "givens" define what is expected from the organization and enable the creation of meaningful "subgoals" or plans by limiting the number of alternatives that the planners themselves must consider.

2.3 The Potential Implications of Algorithms for Organization Charts

These theories highlight that organizational structures have historically been defined by individual human information processing capacity (based on individuals' existing technology use). With increased information processing capacity, the decision domain of individual jobs may shift and impact organizational structures.

In general, scholars have predicted and found that changes in information processing and communications technology do change organizational structures, paying particular attention to effects on decentralization [e.g., 1, 10, 46, 47, 56, 81]. The evolution of digital communications and technologies also enables novel organizational structures, such as flash teams and flash organizations—temporary crowdsourced organizations complete with roles and hierarchies [59, 72]. Some theorists have predicted that current technological trends, in particular algorithms, are likely to impact decision-making and lead to changes in organizational structures [e.g., 76]. Yet researchers have not yet explored how these changes unfold or how the resulting tensions might be resolved. Moreover, little attention has been paid to how the current technological trends of "Big Data" and "Machine Learning" are affecting the processes that produce organizational structures and the quantified performance measures that accompany organization charts and accomplish control and coordination [24, 44, 81].

In sum, many theories suggest that algorithms will have implications for traditional organization charts, but to date, most literature on algorithms in the workplace has focused on algorithms affecting individual users' work practices without following implications for the broader organizational structure. New research is needed to explore why and how algorithms may come into tension with existing organizational structures and how these tensions can be resolved.

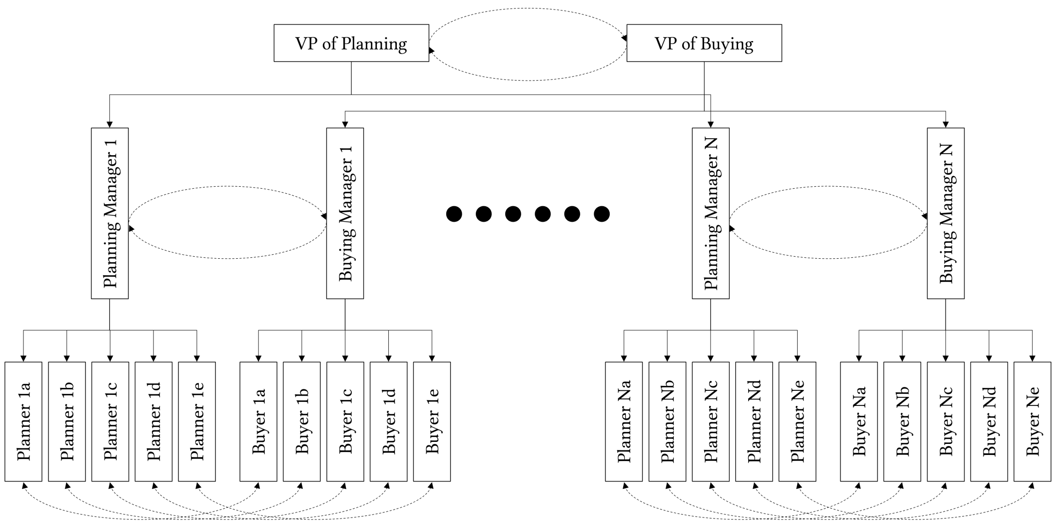

Fig. 2. Reporting structure of the Merchandising Department by title: The buying and planning teams were sister organizations with parallel structures and paired roles at each level of the organizational hierarchy. The titles of roles, their hierarchies, and dotted line relationships are discussed herein; this figure may serve as a useful reference.

3 METHODS AND ANALYSIS

This study reports results from an ethnographic study conducted at a large online retailer, pseudonymously named AlgoCo. AlgoCo had a stated strategy of developing and using proprietary data and algorithms in all parts of the company. AlgoCo had a centralized Algorithms Department, which employed over 100 data scientists and had deployed many algorithmic tools across many functions in the organization. We selected AlgoCo as the context for this study because of its track record of successfully deploying algorithms as our broader research goals centered around understanding the process and impacts of successful adoption processes. As will be discussed, the selection of an organization accustomed to developing and deploying algorithms likely facilitated the study of an algorithm's effects on organizational structures in addition to individuals' work.

3.1 Research Setting: The AlgoCo Merchandising Department

To contextualize the findings in this work, it is useful to understand the Merchandising Department, the purpose of the algorithm developed as well as the Algorithm team's philosophy and process of development.

The fashion buyers of AlgoCo sat within the Merchandising Department. This department contained two functions relevant to inventory assortment planning: buyers and planners. The planning and buying teams were parallel organizations that worked together closely. For each position in the planning team organization chart, there was a paired role in the buying organization chart. Each team of buyers was led by a buying team manager. For large departments, such as the Women's department, several buying team managers might also report to a buying team director. We will refer to these individuals as managers, regardless of whether their title was manager or director. Buying team managers and planning managers reported to their respective vice presidents, as shown in Figure 2.

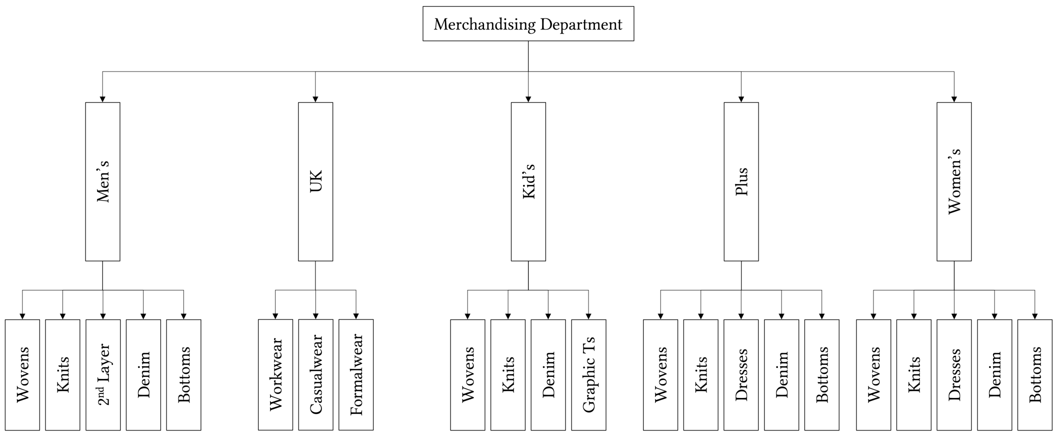

Fig. 3. Structure of the Merchandising Department by Category: Each buying and planning manager was assigned a department for which they conducted inventory planning and purchasing activities. For example, one department was charged with buying Men's clothing. The exact organization of a department by type of clothing varied, depending on the needs of the department. Note, for example, the differences between the organization of the UK and Women's departments. The original delineation of these departments and organization by types of clothing (e.g., Bottoms) are discussed herein; this figure may serve as a useful reference.

At the beginning of this study, buyers and planners were organized into departments and subsequently teams by department (e.g., Men's, Women's) and then type of clothing (e.g., Dresses, Bottoms, Casualwear, Formalwear). See Figure 3. The US and UK teams differed in exactly how these departments were delineated and divided into teams, but the general format was similar; this will be discussed further in the next section.

Roles were delineated such that the buying team was responsible for the high-level assortment strategy and actual purchases while the planning team was responsible for helping to understand how the assortment plan would impact the organization's metrics, such as revenue and margins. Given this delineation, the planning team set performance targets for the buyers.

3.2 Inventory Assortment Planning

Inventory assortment planning is the process that buyers undertake each season to determine the inventory that AlgoCo offers to customers. Previously, the buyers had done this work collaboratively with the planners by using Excel spreadsheets to track their selected styles and calculate the projected metrics for their lists of selected styles (e.g., Brand A's dark wash skinny jeans, Brand B's light wash bootcut jeans) and the "depths" of each style they planned to buy (e.g., 3,000 pairs of style 1, 5,000 pairs of style 2). A mock-up of such a plan can be found in Figure 4. In this example plan, different types of apparel, which would be purchased by different buyers, appear in the rows, and different styles of these clothes, which represent a particular buyer's inventory plan, appear as items in that row.

The success of an assortment plan was determined by several established metrics, including revenue, margins, and "keep rate" (KR), associated with a particular plan. Keep rate was an important metric for AlgoCo as an online retailer that sent customers items based on their personalized style. Keep rate was calculated as the number of customers that purchased an item divided by the number of customers that were sent that particular item. During the rest of the season, buyers would work to secure orders based on this assortment plan.

Fig. 4. Mockup of Inventory Assortment Plan: A department's inventory assortment plan would consist of varying styles of select apparel types. A single buyer was responsible for determining the different styles to stock for one apparel type (e.g., dresses). This figure is an illustrative visualization of such a plan; different apparel types are represented as rows and various styles in each row represent a buyer's inventory plan.

3.3 The Algorithm Team's Approach

The Inventory Assortment Planning Algorithm consisted of both a mathematical optimization model as well as a user interface. The algorithm would recommend an entire inventory plan for each individual buyer based on their buys (i.e., the number of items they needed to purchase for their segment of the inventory) and their constraints (e.g., what percent of their plan should be provided by different vendors, what percent of their plan should include different silhouettes such as sleeveless or short-sleeve). The output would be displayed as a list of recommended styles and the recommended "depths" or volumes to purchase of each style. The screen also displayed all of the calculated metrics for each recommended plan. The data scientists created a visualization feature for the buyers so that they could visualize all of the various recommended plans. This feature helped the buyers understand, evaluate, and choose among the recommended plans.

The data science team took a human-centered approach to the tool development, first observing the buyers' work and then, engaging collaboratively with them to understand their needs and mental models of the inventory assortment planning work. Importantly, the data scientists were agnostic about the decisions of the planners, buyers, and their managers. The data scientists were more focused on teaching these groups how to make and evaluate their own decisions. As a result of this collaborative development approach, and the data scientists' approach of letting the buyers continue to define and control their own decisions when using the algorithm, we observed little resistance from the buyers. This contrasts with prior research which has shown experts' resistance when new algorithms seem to threaten their autonomy or identity [19, 37].

3.4 Data Collection

We negotiated access to study the development and implementation of a new algorithmic system for planning inventory. We chose an inductive field-based research design to match the early stage of the research literature and the developing phenomena [20]. To gather this data, the first author of the current study arranged to work as an unpaid program manager (PM) within AlgoCo's Algorithms Department. As previously stated, AlgoCo had a stated strategy around algorithms, which made it a potentially rich site for study. This arrangement allowed for more access to information about the algorithms, their development, and their impacts on the organization than could be gleaned from public information or understood from other methods, such as interviews. Aside from the stated strategy around algorithms, the author had little prior knowledge of the inner workings of this organization.

She attended the algorithm development meetings with the working team and their managers and executives. She also identified and embedded herself in a specific data science team which had a specific capability under development. We chose to focus on a single algorithm project so that we could study the before, during, and after phases of the development process. Though the author actively participated in the organization as a program manager, she did not participate in the technical aspects of the algorithm's development or in the development of strategies around selecting potential algorithms to develop.

Our research design was inductive; at the time we began the study we did not anticipate that the new algorithmic development project would have implications for the Merchandising Department's organization chart. This finding was emergent; as the study progressed, we began reading and iterating between our observations and relevant research literature.

As an unpaid PM, the first author regularly worked at the company headquarters, located in a large US city. The first author embedded with a specific buying-planning team to study their work processes before, during, and after the algorithm was developed. Through this position, she also interacted with the buying and planning managers and executives throughout the project both in meetings and in regular reflection interviews. She was subscribed to the internal communications, data storage, and knowledge-sharing platforms used by employees. The author was also involved in both formal and informal onboarding and social activities. The author attended team meetings, managers' meetings, and directors' meetings in the Algorithms Department. She also observed user testing meetings, cross-functional governance meetings, or, once the tool was developed, user training sessions. Archival data on the Merchandising departmental structure and organization charts since AlgoCo's founding was also collected.

The study took place over a 10-month period. At the end of the first study period, the Inventory Assortment Planning Algorithm had been broadly adopted across the buyers' organization. The adoption and use were tracked within a dashboard on the tool's landing page. As the first author was leaving the field, the leaders of the Merchandising Department and senior executives were discussing whether and how to change their approach to inventory planning based on tensions that had emerged. The first author negotiated to return to the company one year later to conduct follow-up interviews and observations to see whether and how adoption of the algorithm had continued and changed. This month of data collection also included many interviews targeted specifically on understanding the tensions between the algorithm and the organizational chart that had emerged during the original data collection and to learn whether and how those tensions had changed. The findings from this period are reported below. During this month, the first author observed the same set of meetings as during the initial period (buyer team meetings, data science team meetings, cross-functional meetings). This month also included observing many instances of the buyers doing their inventory planning using the algorithm independently (i.e., no data scientists present), which was a new data source. The tensions discussed and relevant at this period were the same as during the original data collection period as evidenced in the observation notes and interviews.

3.5 Data Analysis

We followed a grounded theory approach when analyzing our data (Glaser and Strauss 1967; Charmaz 2014). As our main argument relates to the significance of tensions and changes over time, our analytical approach was structured to characterize, substantiate, and illustrate these changes. This analysis involved reading field notes, interview transcripts, memos, and archival data several times and coding our data in NVivo.

The first author collected the data and also conducted the first full pass of data analysis, coding each piece of data in an open-ended and inductive process. In collaboration with the third author, they made a key interpretive move which was to focus this particular project specifically on the tension between the algorithm and the organizational chart. Other themes that they discussed but left out of this paper (for analytical clarity) included the buyers' learning and reskilling process to be able to use the new algorithm and the data scientists' human-centered development process. Both of those processes were important for the ultimate adoption of the algorithm. We did not include them in this paper so we could focus on the focal research question about the algorithm and the organizational structure. We chose to focus on this theme because it was well-represented in observations and interviews across all study phases and offers novel theoretical insight to the literature.

The first author coded every piece of data and created a spreadsheet analyzing every piece of data for themes identified by the first and third authors. We saw that within the first phase, before the algorithm was developed, people had a taken-for-granted way of making sense of their decisions, jobs, metrics, and the Merchandising Department's organization chart. Many of those taken-for-granted assumptions about "the way things worked" became visible during discussions about changes to the organization chart (as described in the Findings Section) and also during the development and prototyping of the Inventory Assortment Planning Algorithm. For later phases as the development of the tool progressed, we also conducted a thorough analysis of the many cross-functional interactions that played out as the data science team developed the tool, in collaboration with the 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 during this period in various meetings and interactions.

Having focused on these themes, the first and third author collaborated on analyzing data excerpts for their meaning and significance within the research question. The first and fourth authors discussed these themes and findings throughout the data collection and analysis process, but the fourth author was not involved in the line-by-line analysis of every piece of data. Instead, the second and fourth authors worked to help theorize the findings that the first and third authors had produced through many rounds of iteration and connect them to the literature.

4 FINDINGS

During our study, a team of AlgoCo data scientists developed a new algorithm that helped the buyers with the work of inventory planning. The algorithm replaced the spreadsheets previously used by buyers and was configured to recommend the styles and depths for buyers to include in their inventory plans. However, as the buyers began to use the new Inventory Assortment Planning Algorithm, they began to see that their old way of producing inventory was coming into conflict with the new and evolving algorithmic approach. Previously, the buyers had used the organization chart to divide inventory decisions into product segments aligned with buying

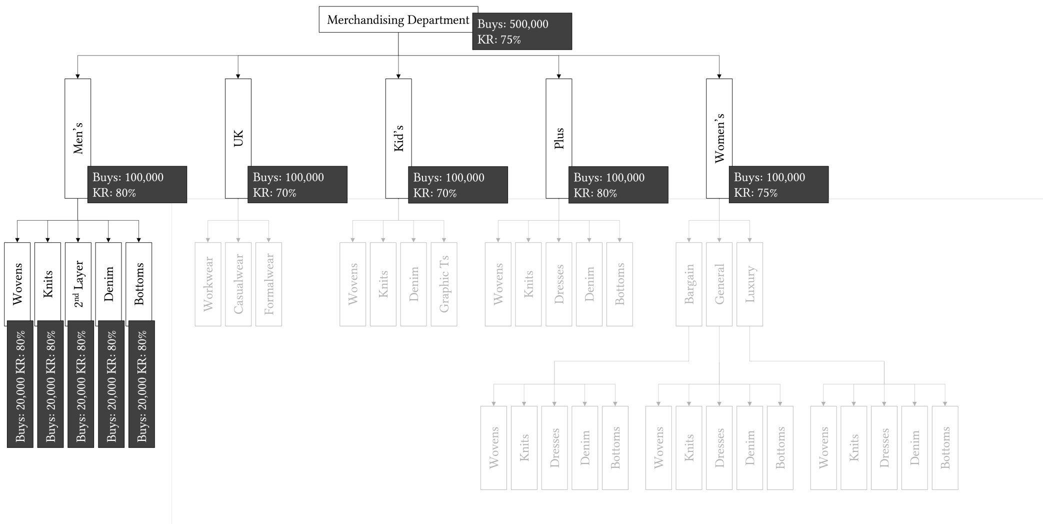

Fig. 5. Organization chart overlaid with subdivided buy targets: The number of buys and the related target metrics such as keep rate are divided between teams such that the number of buys and metrics at the level of an individual buyer aggregate to the target number of buyers and metric performance at the department level. Department level targets then aggregate to the targets for the whole Merchandising Department. This concept is illustrated with a set of example buy and keep rate targets for the Men's Department.

teams and sub-segments aligned with roles on each team, but they soon realized issues with this segmentation approach. The data scientists and buyers saw that the algorithm produced different inventory plans based on the decision space it was given (e.g., recommending plans for the entire buying team or each individual buyer) and that they could compare the respective performance of those plans. When the algorithm could explore a larger decision space, it could recommend inventory plans that performed better on established inventory plan metrics, in part because it could identify styles missed by the segmentation approach and complementary styles previously isolated in different roles, and it could suggest more flexible changes in style distributions.

To develop these ideas, we report findings that show that the algorithm and the organization chart were both being used to help organize and coordinate a large set of decisions, but that they mobilized different approaches to that problem. In this findings section, we analyze the difference between these two approaches and why they were in tension.

4.1 The Organization Chart Coordinated the Set of Decisions Involved in Inventory Planning (Baseline; Month 1)

The buying and planning managers used the organization chart to subdivide a large set of decisions into smaller domains (buying teams) and then even smaller domains (individual buyers' jobs). For example, if they needed to buy 500,000 items into the inventory for an upcoming quarter, they would subdivide those 500,000 decisions into 5 teams: 100,000 each for Women's, Men's, Plus, Kid's, and UK respectively. They would then further subdivide the 100,000 units assigned to the Men's Buying team among the 5 buyers on the team, assigning each of them 20,000 items of inventory to decide which styles to stock, and at what depths. See Figure 5.

The buying and planning managers then also used the organization chart to control and coordinate the performance of each of these smaller decision domains. They would assign each buying team and each buyer a set of performance targets that their inventory plan needed to hit each quarter. For example, the Men's Denim Buyer would be assigned 20,000 "buys" and would be given targets related to the profit margin, revenue, and customer satisfaction that she was expected to hit with her inventory plan.

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). The process of dividing up the decisions and assigning targets was accomplished in collaboration with the Finance Department. The process was owned primarily by the Vice President of Planning (the top position in the organization chart shown in Figure 2) and accomplished in collaboration with the planning managers (who were each paired with a buying team manager). The process was informed by historical data from prior quarters but was also fairly "manual," in the sense that the planning manager would divvy up units across teams and then iteratively move those around as she balanced inventory across teams. Note, this approach is similar to most retail companies and has been used for over a hundred years.

4.2 The Organization Chart is Designed to Structure a (Human-) Manageable Set of Decisions (Baseline; Month 1)

In our study, we observed several instances where the Merchandising Department's organization chart changed. These instances illustrate how the managers and employees were using the organization chart and its assumed purpose. The discussions around these changes illustrate how the organization chart 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. The Plus Buying Manager explained the decision:

> We split out tops into someone who was responsible for wovens and someone who had responsibility for knits and sweaters, just to make the scope of responsibility more equitable and more manageable. (Buying Manager 3)

The decisions and targets for the Tops Buyer were thus segmented into two buyer roles—one buyer was responsible for developing the inventory for Wovens, while another was responsible for Knits and Sweaters. Each buyer was assigned their own volume and targets. There was no discussion of whether this division would impact decisions, targets, or outcomes; it was an assumed, taken-for-granted division of labor based on the growing sales and the need to split the number of decisions for human manageability.

The second example involved the company newly entering a new market in the UK. The Merchandising Department expanded to include a UK Buying team alongside the Women's, Men's, Plus, and Kid's Buying teams. The UK executives who formed and structured the buying team decided to structure the buying teams based on how the customers might use the clothes, rather than by product type, the more standard structure. The UK Buying team thus had an Eveningwear buyer, a Casualwear buyer, and a Workwear Buyer (instead of a Wovens, Knits, and Denim buyer, as on the other buying teams). As the UK Buying Department was being structured, this non-standardized way of structuring the buying roles was easily accepted by the Merchandising

Department and AlgoCo executives. It was explained as the way of structuring and dividing out the decisions that was most manageable and useful for the UK buyers. Later, those non-standard roles and product categories introduced complications for some of the data science approaches, but with the traditional way of understanding the organization chart, this division was straightforward. There was no discussion of whether dividing out the decisions by customer end-use would impact the decisions, targets, or outcomes.

4.3 A New Algorithm Recommends How to Make Sets of Decisions (Month 4)

The sections above explain how the Merchandising Department was using its organization chart to help organize and coordinate the large volume of decisions that had to be made each quarter to produce their inventory. During our study, a team of data scientists created a new algorithm that came into tension with this method of producing inventory. The new algorithmic approach began to show that using the organization chart to divide out the decisions was inventory plans for the department that buyers viewed as poorer performing, in ways the Merchandising Department had not ever realized or considered.

4.3.1 Data Scientists See Organization Chart as Decision Tree and Design for the "Leaf Nodes". Data scientists' conceptualization of the Merchandising organization chart as a decision tree started to reveal how the organization chart was affecting the design of inventory plans. Several data scientists in various meetings talked in offhand ways - meaning most people there understood the point - about how the Merchandising Department organization chart (recall Figure 3) and its parallel data structure could be understood as decision trees. The data scientists 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 a data interface that organized all the items in the AlgoCo inventory. He used "earrings" as an example product category:

> See how earrings has a parent in the tree (i.e., jewelry, the category it is nested in) and jewelry has a parent in the tree (i.e, accessories, the category jewelry is nested in). There are some things that if you follow down, nothing has them as a parent. Those are leaf nodes. (Data Scientist 10)

He then emphasized, "So those (gesturing to a 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 to design algorithms for workflows that have to involve leaf nodes and these bottoms up decisions. (Data Scientist 1)

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."

4.3.2 Buyers Curate Sets of Algorithmic Recommendations in "Leaf Nodes". The data scientists collaborated with the buyers to develop a new inventory assortment planning algorithmic system. In a short time, the data scientists were able to model inventory planning as an optimization problem, where the front-line ("leaf node") buyers were making decisions about inventory to stock in ways that optimized the performance of their segment. Before the development of the algorithm, buyers used spreadsheets to calculate the projected metrics of a set of styles for an inventory plan. Buyers would iterate between their plans and their projected metrics through somewhat manual and tedious calculations in the spreadsheets. Given their buys (e.g., 20,000 items of denim), they were picking several dozens of styles and depths (e.g., 1,000 of Style A, 1,000 of Style B, 1,000 of Style C) and needing to calculate the projected metrics for the decisions.

The data scientists implemented a human-centered algorithm design process to develop the new Inventory Assortment Planning Algorithm. They shadowed and interviewed the buyers and learned that they could reframe the buyers' inventory planning process as an optimization process. The mathematical model that they developed for the inventory planning was in the form $Z = \mathrm{a}x + \mathrm{b}y$, where Z is the metric being optimized and x, y are the decision variables. For example, a buyer might want to optimize for keep rate. In this case, she would specify keep rate for Z. She would then stipulate other conditions for x and y. In this case, x could be the percentage of items that must be red, and y could be the minimum number of styles in the assortment plan. The data scientists worked closely with the buyers to determine the metric to optimize and the input variables to include in the model. The data scientists engaged the buyers' help in figuring out the metric to optimize, and what input variables to include in the model.

The buyers would input their buys (e.g., 20,000 denim items) and constraints (e.g., 10% red) and then push a button for "get recommendations." The algorithmic output was a recommended inventory plan. See Figure 6. The algorithm would display a list of recommended styles and depths that would optimize the specified variable, conditional on the inputted constraints. The left side of the screen displayed all of the calculated metrics for that potential plan. The algorithm calculated these recommendations and projected metrics within seconds, whereas the plans and calculations used to take weeks to calculate. The planners would do many calculations "by hand" in Excel, meaning adjusting numerical values and using macros to calculate impacts cell by cell and sheet by sheet. The data scientists added many features to help the buyers understand the algorithm and its output, including a visualization screen where they could see all the styles they had picked and arrange them in a pivot table across many different dimensions.

This new algorithmic Inventory Assortment Planning Tool introduced buyers to insights and practices that influenced their perception of how the organization chart was shaping and constraining inventory recommendations. First, as the buyers, buying, and planning managers began to use the tool, they learned how to more systematically measure the impact of different decisions on various outcomes. As an example, previously, buyers had not used client segmentation variables in their planning. Inclusion of client segmentation variables might look something like knowing that 30% of projected customers were going to be over age 50 and choosing styles based on those clients' preferences. By providing automatically calculated metrics and inventory plan visualizations, the Inventory Assortment Planning Tool made it easier for the buyers to compare the projected performance of an inventory plan that included client segment variables against an inventory plan that did not include the client segment variables. Within minutes, buyers could see that their keep rate might be materially higher if they implemented the recommended inventory plans which leveraged customer segmentation variables.

Second, as buyers learned to use the tool, they also learned how imposing constraints impacted the projected performance of their inventory plans. Such constraints might be strategic and/or necessary (e.g., remove an out-of-business vendor) to make an algorithm's recommendation viable

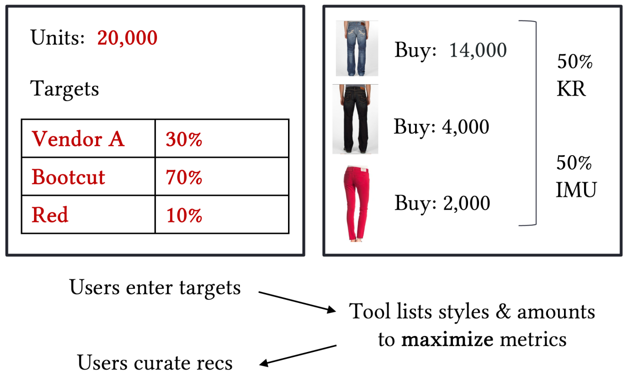

Fig. 6. Mockup of Inventory Assortment Planning Tool User Interface: The Inventory Assortment Planning Tool allowed buyers to input various targets and constraints, as illustrated by the box on the level. Buyers received recommendations for inventory plans which optimized for the input target given the set of constraints. Buyers were also provided a set of projected metrics, such as keep rate, for the recommended inventory plan. These outputs are shown in the box on the right.

[60]. However, the buyers also learned that they were imposing constraints based on their intuition, such as limiting the resulting inventory to a certain color, in ways that impacted the performance of their inventory. The Inventory Assortment Planning Tool allowed the buyers to now better understand the impact of constraints on the projected performance of their inventory plans. As an example, buyers now could see that if they included a constraint that inventory plans needed to include 10% red styles, the resulting plans produced by the Inventory Assortment Planning tool had a lower projected performance than if the constraints were not included. The data scientists did not argue about the appropriateness of any constraints. They simply taught the buyers to measure the impact of those constraints themselves. As an example, one of the data scientists regularly said during prototyping sessions where the buyers were learning to use the Inventory Assortment Planning Tool, "Let's put it in and see" about different constraints. The buyers would input the constraints and learn for themselves how those constraints were influencing the recommendations. The data scientists conducted many ongoing training sessions for the buyers and the buying and planning managers about optimization, constraints, and trade-offs using the new system.

4.4 Data Scientists and Managers Discover Issues Between the Algorithm and the Organization Chart (Months 5-9)

The data scientists configured the new algorithm for the front-line buyers to use to get recommended inventory plans for their individual roles and decision domains. Recall from Section 4.1 that buyers' jobs and related decision domains - i.e., "plan the 20,000 items of denim for the next quarter" - were not structured in a standard way (e.g., differences between the Women's Buyers' jobs and the UK Buyers' jobs) and were designed to support manageability and practical decision-making for one human person (e.g., splitting tops into Knits and Wovens as the volume in that category grew). As the buying teams learned to use the new algorithm, they began to see how the organization chart was constraining the algorithm's recommendations and resulting in poorer-performing inventory plans, given a focal and established metric for inventory plans.

4.4.1 The Organization Chart Was Arbitrarily Segmenting Decision Space. The data scientists had chosen to design the Inventory Assortment Planning Tool for use 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 and 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 (Data Scientist 16) in an interview:

> **Interviewer:** OK so what you all are saying is… Consider two scenarios. In the first you set up two buyers' roles like the UK Buying team did: 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 like the Women's Buying team did: 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 organization chart itself was segmenting decisions into jobs through the assignment of buys and targets and that this segmentation affected resulting inventories and subsequent metrics, like keep rate. Specifically, the organization chart was segmenting decisions into human-interpretable decision spaces based on recognized product taxonomies that were simply taken-for-granted. Experimentation with the new algorithm, however, called these divisions into question. When buyers and the data science team ran the algorithm without the constraints imposed by the organization chart, the recommendations were predicted to have better outcome metrics (e.g., margin, revenue, keep rate) than the aggregate of metrics associated with inventory recommended for each buyer individually.

Note, this issue offers an important example of how the algorithm design choices deeply impacted the decisions generated by model. Many prior studies, for example, Suresh and Guttag [69], Bucher [12], and Lustig et al. [42], similarly show how algorithm design choices can impact the outputs in socially constructed and often arbitrary ways. These cases matter because they highlight how algorithms influence how we view and structure the world, but are often driven by somewhat subjective and arbitrary decisions. And specific to our research focus, these different recommendation sets (i.e., those produced by humans structured in the org chart as compared to those produced by the algorithm) also revealed that the organization chart divisions had been enabling the buyers to ignore interdependencies between their product categories, which the algorithms' recommendations later surfaced [e.g., 67, p. 241].

4.4.2 The Organization Chart Was Constraining the Exploration of Other Optimal Solutions. Both groups came to recognize the segmentation as a problem and discussed it in meetings and interviews. A data scientist expressed the problem this way: "We saw the algorithm could explore a larger space for better results" (Data Scientist 8). A buying manager built on this idea. She said, "Segmenting our teams and inventory in this way [the org chart] 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." One of the executives said in a strategic planning meeting, "I fully understand the drawbacks of how we are currently organizing the Merchandising Department. We are just now figuring out the better way" (Executive 4). 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" (Data Scientist 16).

4.4.3 The Organization Chart Structures Only One Dimension for Decision-making. The buying manager's quote above also highlighted that the organization chart was not letting them plan inventory in a multi-dimensional way. As one example, they struggled to integrate customer variables and insights. Relatedly, they saw that segmenting decisions by men's and women's or luxury and general products meant they missed certain styles (e.g., androgynous styles for non-binary customers or styles that fell between luxury and general). The Vice President of Buying explained a similar reflection. She said that by segmenting out the buyers' jobs based on a product taxonomy (recall Figure 3), the buyers had to buy products with the average AlgoCo customer in mind. She elaborated:

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

The buying managers had the sense that their inventory plans would perform better on their desired metrics, such as keep rate, if they were able to include more data and insights on clients in an upcoming quarter. In fact, before the new Inventory Assortment Planning Algorithm was developed and implemented, the Merchandising Department attempted to use a moment of departmental restructuring (which they called a "re-org" or reorganization) to bring more client insights into the inventory planning process.

There were many conversations about how to use the re-org to have the buyers focus more on specific customers. For example, at a multi-day "off-site", executives, buying managers, and data scientists all discussed how to split up the Women's Department. The data scientists from the Algorithms Department wanted to divide up the Women's 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 keep rate. They argued for structuring the Merchandising Department based on attributes most 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 data science team had shown that Age was the client attribute that most strongly conditioned keep rate. (Data Scientist 8)

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. One of their executives explained that she did not think that customers' preferences were that different based on their ages (the data scientist's proposal), so she thought developing inventory targeted to the ages would not produce inventory that performed better on their established inventory metrics. In the end, the buyers' authority for their own department prevailed, and the Women's Department was divided up into Bargain, General, and Luxury price points. This re-org was responsible for an additional layer in the Women's Department hierarchy which was another non-standard structure and segmentation of the decision space. This example illustrates how the data scientists and buying group both struggled to include client insights in the inventory planning process as structured by the organization chart.

4.4.4 The Organization Chart Was Modeling Relationships As Hierarchical That Were Not Hierarchical.

One of the data scientists recognized a related problem, which was that they had modeled their merchandising data structure after the organization chart. He led the work of conceptualizing why it was a problem to have the data structure be hierarchical like the organization chart, rather than "flat," using labels that did not nest into a hierarchy. He wanted to convince other data scientists and AlgoCo leaders to work on decoupling the data structure from the organization chart's "people structure." 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 in the Algorithms Department 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." (Data Scientist 16)

His description was explained that the data structure was in fact encoding the people structure, rather than recording the properties of the items themselves. To say it a different way, if anyone looked at 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. He said in a presentation 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 (i.e., the people structure or the org chart).

4.5 Data Scientists and Buying Managers Work to Resolve These Issues (Month 9 and Year Follow-up)

The data scientists and buying managers had thus discovered ways that their previous approach to inventory planning using the organization chart to structure individual buyers' decision domains and performance targets was at odds with the new algorithmic approach to inventory planning. At the end of our study period, the individual buyers were using the algorithm to create algorithmically recommended (and buyer-curated) inventory plans in their "leaf nodes." However, having recognized these problems, the AlgoCo managers also worked to innovate solutions to the tensions between the old ways of using the organization chart to produce inventory plans and the new algorithmic capability to do so.

4.5.1 Let the Algorithm Explore Solutions: "Roll up the Leaf Nodes".

The data science team proposed a solution for letting the algorithm explore a larger solution space, rather than one that was segmented by individual buyers' job domains. Recall that the data scientists saw the organization chart as a decision tree and considered the structuring of the leaf nodes (the buyers' jobs, where the decisions were made) to be arbitrary and to be unnecessarily constraining the search space. One phrase that caught on referred to the idea to "roll up a leaf node" and run the optimization recommendation algorithm there. Figure 8 visualizes what was meant by this idea. The original

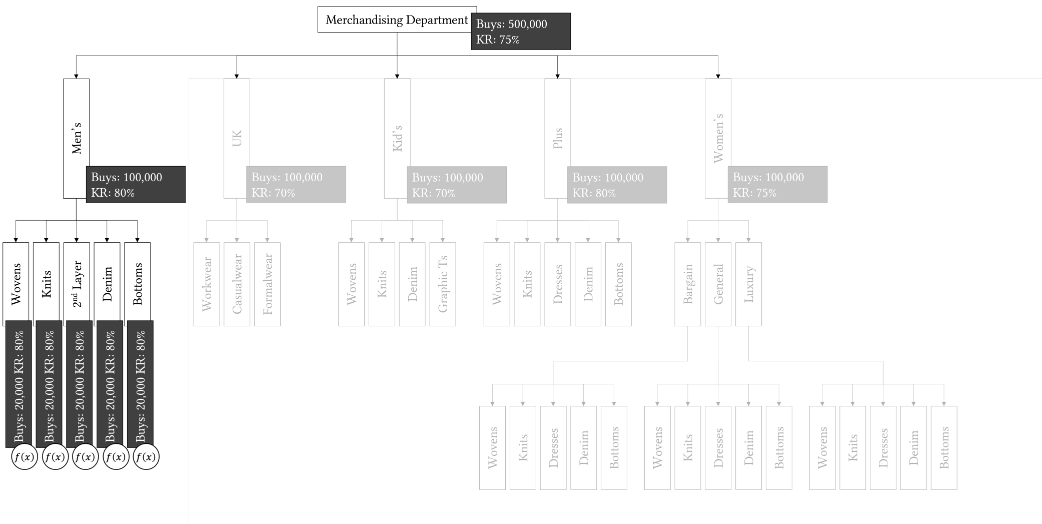

Fig. 7. Original location of Inventory Assortment Planning Algorithm relative to the Merchandising Department's buying targets: Revisiting the example targets in the Men's Department shown in Figure 5, buy targets and metrics were divided between buyers. Originally, the Inventory Assortment Planning Tool operated at this individual buyer level, as illustrated by the circles containing $f(x)$ shown.

configuration of the algorithmic tool was to produce an optimized set of recommendations for 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. This idea was the specific way of allowing the algorithm to "explore a larger space for better results."

Specifically, the data scientists proposed experimenting with "rolling up the leaf node" on the Plus Buying team. This proposal meant that the manager and leaf node structure typical of organization charts 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. This process would eliminate the need to buy targets at the individual buyer level. The Plus Buying Manager was willing to try this experiment and learn from the process of group-level algorithmic recommendations and curation. She and the data scientists envisioned that the buyers on the team would take on more flexible roles that might change every season rather than being persistent and defined by product types (e.g., Plus denim buyer). They brainstormed having the more flexible roles change in response to the algorithmic recommendations (e.g., perhaps having a dedicated denim buyer if the algorithm recommended more denim one season) or be specialized along other dimensions (e.g., specializing in vendor relationships). Both the data scientists and the merchandising executives recognized that these changes would alter the manager-buyer relationship (especially in terms of accountability for decision-making) and the overall role structure of the buying team. They agreed on an experimental approach where they would try different configurations and learn from them over time.

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

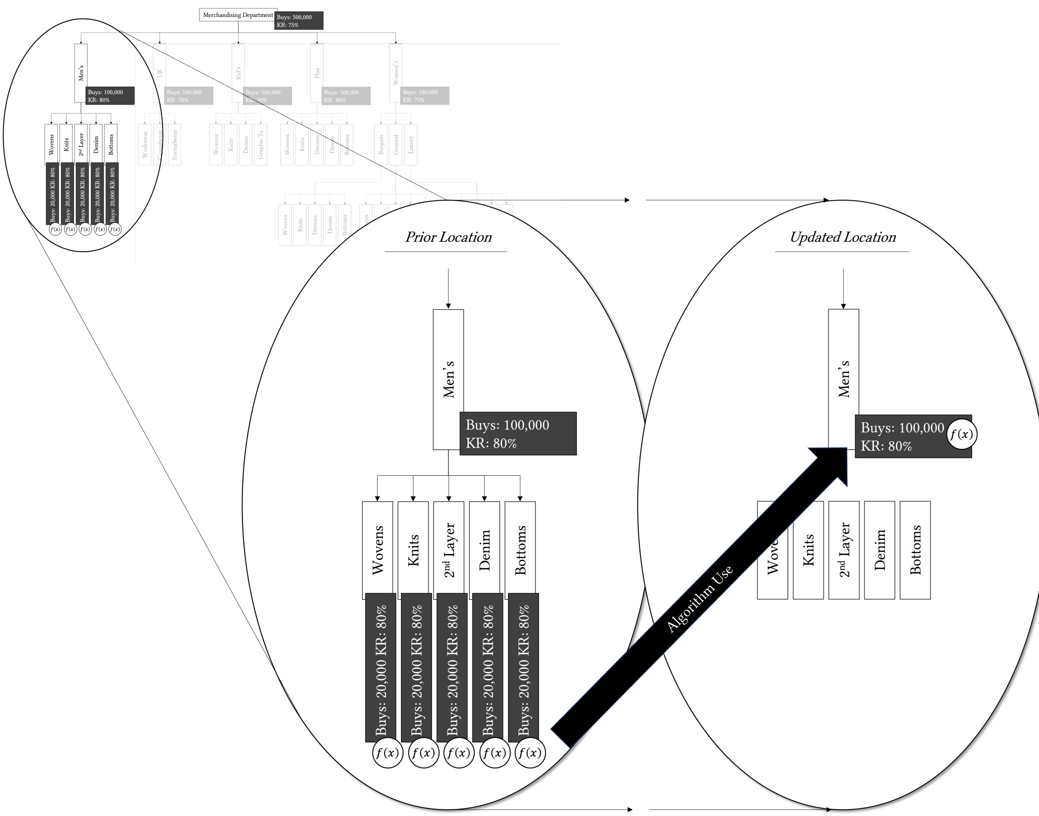

Fig. 8. Updated location of Inventory Assortment Planning Algorithm relative to the Merchandising Department's buying targets: The Merchandising Department and the data scientists who developed the Inventory Assortment Planning Tool proposed allowing the Inventory Assortment Planning Tool to explore the decision space of the team, rather than being constrained by the decision space of a single buyer. Using the same example of the Men's Department, as shown in Figures 5 and 7, the Inventory Assortment Planning Tool would be used at the Men's Department level, as indicated by the circle containing $f(x)$. This meant that buy targets and metrics would no longer be used at the level of the individual buyer, but remain at the departmental level as well, as illustrated.

the inventory planning decisions could be structured. As an example, the data science team kept on their team roadmap 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, et cetera." One of the data scientist's strategic ideas 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 (Data Scientist 1).

These continuing discussions at AlgoCo about how they reconfigure the people structure, the data structure, and the data tools are well-summarized by some educational materials that one of the data scientists put together, which highlights how the data labeling can capture fine granularity, such as a combination of silhouette and class, but can be rolled up to higher-level decision spaces:

> 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.
> 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 both "people structures" that are practical and interpretable for people's decision-making and algorithmic approaches to dividing and aggregating decisions and outcomes. Their aim was to flexibly balance these approaches going forward.

4.5.2 Replace Hierarchical Data Structures With Flat Data Structures. AlgoCo started to make changes to pursue balancing these two approaches. One change was to reorganize the data structure to eliminate the hierarchical structure patterned after the Merchandising Department's "people structure" or organization chart. One of the data scientists proposed a solution - he argued that it was much more consistent with data science approaches to store the data "flat" and use flexible "roll-ups" instead of static hierarchical divisions. He explained:

> The proposal is to build a new data model where we use hierarchy only for concepts that are truly hierarchical. When a hierarchy is not unambiguous, tags are better.
> My example here is, in old email clients there would be folders, and you would have folders within folders. That is 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 choices about where it goes; you just tag it with every tag that's relevant. (Data Scientist 16)

His point was that by structuring the data structure following the organization chart, their systems were losing items in unnecessarily hierarchical structures. The hierarchical data structure would store information on one "women's woven top" only in that segment of the hierarchy. In contrast, if they decoupled their data structure from the org chart, and made it flat, that same item could be tagged with as many tags as possible and then could be flexibly seen, included, and rolled up into sets such as "any green item" or "anything from Vendor A" or "anything for millennial clients". His vision was for every item to be tagged with as many relevant tags as desired and then "roll-ups" could flexibly aggregate relevant items using tags depending on a focal analysis. Modeling the data structure after the org chart had prevented this functionality.

4.5.3 Work to Decouple People Structures and Data Structures. The plan to change the data structure to a flatter and more flexible model where all items were tagged rather than stored in hierarchies was a huge undertaking, but was also well-received within both the Algorithms team and Merchandising. One of the buying managers explained it this way in a meeting, "We are thinking about breaking the dependency of the data structure hierarchy and how buying 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 managers saw opportunities in terms of how the Merchandising teams were staffed and structured.

AlgoCo also worked to reconfigure the buyers' organizational hierarchy as they came to see the issues that were created by the way the organization 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 organization charts that did not change very often. They had come to understand that the organization chart 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. Managers in both departments saw strategic opportunities to separate out these different structures and more flexibly and dynamically model some of the decisions that were being constrained by the static organization chart.

4.5.4 Short-term Implications for Buyers' Collaborators. The focal finding of our paper was this tension between the new algorithm and the buyers' organization chart. We note one final implication of the algorithm related to the planners. The data scientists' project was to create an inventory planning tool for the buyers. In so doing, they ended up automating many of the manual computation tasks that were involved in inventory planning—tasks that had previously been done by the planners. The buyers at the "leaf nodes" were producing inventory plans using the algorithm in meetings where no planners were present by the end of our study. Some planning managers were involved in the design of the "roll up the leaf node" experiments. Despite some of the planners' tasks being automated by the algorithm, in our interviews with planners, they described feeling busier than ever. One planner told us about the task force she was on related to the data model and data attribution work. She explained, "I'm on a lot of the kind of technical work that we're doing to help first get better organized with our data." More research is needed to explore how these changes develop over time, but in sum, at the end of our study at AlgoCo the buyers were aware of and working out the tension between the new algorithm and their organizational chart, while many of the planners' manual computation tasks had been automated, and they were moving into new task domains.

4.5.5 Pending Organizational and Occupational Change. In sum, throughout our study, managers and executives in both the Merchandising Department and data science team discussed this tension and possible ways to organize the buyers to make use of the algorithm's recommendations. They recognized this as a challenging issue that required new ways of thinking as well as extensive organizational and occupational change. When the first author returned to the field after a year, they had made some progress towards thinking about different responses to this issue. As described, one team was willing to try the team-level recommendations and planning. Interestingly, one of the data scientists had left the company to oversee the inventory of another company that was going to do a different way of solving this tension, where they did department-level algorithmic planning and recommendations.

By the end of the study, the changes at AlgoCo had not yet coalesced into a full re-organization of the buyers' department or new performance expectations for the buyers. Although we had an extensive period of observation, with enough depth and embeddedness to be able to deeply characterize this tension which is missing from the academic literature, we expect that level of organizational and occupational change might take many more years. Whether and why the organization might continue an evolution where its organization chart grows to mimic the algorithm's boundaries, or whether and why it might land in an intermediate compromise solution that balances structures that support human and algorithmic decision-making, remains to be seen.

5 DISCUSSION

Through our 10-month ethnographic study of AlgoCo's development and adoption of the Inventory Assortment Planning Algorithm, we show how a tool built to help individual buyers changed not only their individual processes, but processes throughout the department and in fact came into tension with the entire department's organization chart. In this section, we review and generalize our findings by discussing research contributions as well as potential implications for designers of algorithm-based tools.

We found four related tensions in how the buyers used the organization chart to structure their decision-making around inventory planning compared with how they and data scientists began to configure an algorithmic approach to inventory planning. First, our data showed that the organization chart was arbitrarily segmenting the decision space, which influenced how and what decisions were made. Second and relatedly, these role-based segmentations of the decision space arbitrarily constrained the algorithm's explorations of "optimal" (meaning produced by an optimization model) inventory plans. Third, the organization chart was also only accommodating one dimension for decision-making by factorizing buyers into roles defined by a product taxonomy. There was a "Denim" buyer, but not a "Styles for Women over 40" buyer. In contrast, the algorithm could be configured to take into account both of those variables, but the buyers struggled to do so using the older approach to inventory planning. Finally, the data scientists had inadvertently modeled some relationships as hierarchical which were not hierarchical because they had developed their data structures around the Merchandising Department's organization chart. The buyers, planners, and data scientists together recognized these issues and worked together to figure out new relationships between what they called the "people structures" and the data structures, including the way the recommendations were produced.

Taken together, these findings contribute to the CSCW literature on algorithms in the workplace and also the CSCW and organizational theory literature on organizations as information processing systems.

5.1 Contributions to research on algorithms for individual decision-making

To date, prior CSCW literature on algorithms at work has focused primarily on individual decision-making. Tension between algorithms and existing work practices is a common theme in such studies of algorithms and individual work, as well as in studies of technology adoption more generally. Such tension arises from design decisions, such as what data is used in algorithms [e.g., 74], how classifications are based on this data [e.g., 55], and how experts are expected to leverage the results of these algorithms [e.g., 18]. For example, Petersen et al. found that caseworkers resisted documenting their practical and situated categorizations of welfare seekers, as they felt that outsiders would not understand the context of these classifications and would create an unintended permanence in a welfare seeker's classification [55]. Many of these tensions arise from the use of algorithms to control work [37, 53] and differences in the objectives or incentives of powerful stakeholder groups and end users for the algorithms[e.g., 51].

Our paper differs from and extends this literature by theorizing how algorithms impact organizational structures, rather than focusing on individual work practices and roles. We showed how the organization chart had been used to structure and control inventory planning for a large department of around 200 people and how the new algorithm, initially configured for individual use, ended up coming into tension with other jobs and the overall organization chart. This finding extends the prior studies of algorithms focused on individual work by connecting that literature to theories of technology change which predict role structures and organizational hierarchies will be impacted by "transformative technologies," alongside individual work practice [3, 7].

A second, related contribution involves our inductive finding that this algorithm, configured to aid the decisions of the individual buyers (in the "leaf nodes") ended up calling into question the work of the planning managers, who set the numerical buys and targets that structured the buyers' work. The planning managers were the ones whose decision-making structured the organization chart into manageable jobs for each buyer and then patterned the buys and targets around the organization chart. They had been doing this work for over a decade at AlgoCo, and the impact of the organization chart on that decision space had never been noticed or called into question. This finding is interesting because, in some ways, it is still focusing on an individual's decision-making - the planning managers were indeed making decisions, but their decisions structured the work of an entire 200-person department. So when the algorithm started to come into tension with those decisions, it was not just the planning managers' own work and domain that was implicated, it was an entire organizational structure that was now called into question. To our knowledge, this kind of effect has been anticipated [e.g., 76] but not yet empirically demonstrated in the CSCW or organizational literatures. To understand the generalizability of these findings, we welcome investigation of other settings introducing algorithms into departments that share similar patterns of managers creating plans that structure individuals' work characteristics .

Our study is inductive and ethnographic, precluding causal claims about mechanisms, but our observations suggest that these changes are related to key dynamics that can be explored in future research. For one, the algorithm offered increased information processing capacity for individual buyers. This mattered because buyers' job domain size had been loosely based on their information processing capacity, specifically the capacity to make a certain number of decisions about a certain number of inventory items. Of course, the algorithm offered this increased capacity only in the planning and evaluating of the inventory; buyers still had to execute those plans by negotiating with vendors, securing the purchased items, and monitoring the performance of the inventory as customers started to purchase items.

Though this study centers around buyers and planners in a retail organization, we suspect that many organization charts similarly divide out decision spaces to small collections of similar decisions and deprioritize interdependencies between these decision spaces. Examples might include segmenting engineering departments around product lines, sales departments around geographies, and client services around industry targets. Future research can explore how individual practices and role structures change as planning and evaluating task demands change, as well as how these changes influence relational and execution task demands.

5.2 Contributions to research on organizations and coordinated decision-making

Our study also contributes to CSCW research on organizations, information processing, and coordinated decision-making [4, 8, 36, 67]. Many studies in this area, and in the related field of organizational theory, have shown that changing information and communication technologies co-evolve with changing organizational structures [e.g., 6, 7, 31, 56, 67]. Our study differs from and extends this prior literature by exploring the changes associated not with information or communication technologies per se, but with a new algorithmic system designed to augment individual decision-making. Prior literature showed that changes to information and communication technologies tended to also involve changes to organizational structures because they would alter communication patterns and therefore relationships between roles and groups [e.g., 33, 36] as well as task interdependencies between different roles and groups [e.g., 32]. In focusing on a different technological change, our study suggests related dynamics for change that can be explored in future research.

One dynamic includes the changing information processing capacity of individual roles. The previous structure of the Merchandising Department "factorized" [67, p. 241] inventory assortment planning decisions, starting with dividing out the large set of decisions involved in planning the whole inventory down into broad buying teams (e.g., Women's, Men's) and then dividing those teams into jobs by clothing type (e.g., Dresses, Bottoms). Few studies have documented changes to information processing and related changes to an organization's current "factorization" structure. Our study suggests that AlgoCo responded to the change in individual information processing capacity by "rolling up the leaf node," meaning moving the set of decisions up a layer in the organizational hierarchy. Instead of each individual buyer taking on a more granularly factorized area, the department explored whether a manager and team could take on the broader set of decisions together. This proposed change required new models for structuring the buyers' roles and role responsibilities because the buyers' roles and role structures had been relatively stable for the past decade.

A related dynamic involved the new and different information that became available for planning managers in how they structured and controlled the inputs and outputs of different roles. Previously, the planning managers would take the structure of "factorization" of the department as a "given," meaning that they would plan for certain individuals to make certain purchases every year. In other words, the planning managers themselves had a human-manageable set of decisions and alternatives related to how to structure the targets and metrics based on the relatively stable set of buyer roles.

The Inventory Assortment Planning Algorithm, however, was not limited by the number of alternatives it could consider when recommending the number of buys of each product, and the related targets and metrics. Recall that the data scientists described how the algorithm could be configured to recommend "arbitrarily many" different ways of modeling the entire team's set of buys, targets, and metrics. These many different recommended plans did not have to be constrained by buyers' roles and could consider other dimensions like forecasted customers. It could also take into account interdependencies at the team level without requiring different team members to actively communicate plans and information to each other. This finding suggests that as the algorithm could both recommend and calculate the associated metrics with each recommended plan for the entire team, it was taking on some of the coordination work that the organization chart (as an information processing structure) had previously been doing. Of course, the implementation of the algorithm did not eliminate the need for certain tasks to be divided up and coordinated; there were still tasks associated with purchasing inventory that had to be factorized to an individual buyer's capacity.

These findings provide evidence of an algorithm having an impact on taken-for-granted organizational processes and institutions. Though previous research has touched on the tension between algorithms and existing organizational processes, the conclusion of this research is often that while algorithms might be intended to disrupt the status quo, proponents of the algorithm still depend on conventional practices and entities to achieve their aims [e.g., 41, 61]. Here we find the opposite: the algorithm was not initially intended to be disruptive, but had unintended consequences for the long-established, taken-for-granted practices for structuring buyer's and planner's work.

Overall, our findings highlight for the academic literature the ongoing tension and interplay between algorithms and organizational charts that are likely playing out in many organizational settings. This case shows how a new algorithm informed the organization chart. Though not foregrounded in our study, the organization charts of both the buyers and the data scientists also informed the algorithms - both in terms of what algorithms were developed and how they changed work. We hope that future research continues to explore the ongoing interplay between algorithms and organization charts.

5.3 Limitations and Boundary Conditions

Though information processing capacity has long been considered a key factor in dividing labor and thus determining organization charts [67], this study of an algorithm's impact on information processing capacity and subsequently an organization chart is limited in its scope to one department within one organization. The Merchandising Department's organization chart has characteristics that may be shared with organization charts in other fields or work functions. Such characteristics include (1) work is segmented into discrete categories each involving similar decisions applied to each category, (2) this segmentation limits consideration of interdependencies between decisions in each category, (3) consideration of such interdependencies could improve performance (as defined by the target/goals set for each category), and (4) a managerial position determines the targets or goals for each of these groups and potentially subgroups. Our findings may generalize to departments with these characteristics, such as engineering departments organized around product lines, sales departments divided around geographies, and client services segmented around industry targets. However, future research is needed to establish that our findings are generalizable outside of the retail sector and the specific work carried out by buyers and planners in the inventory planning process. Additionally, given that this study was an ethnography, future work is also required to establish causality between the algorithm, information processing capacity, and ensuing organizational changes.

Additionally, as described, AlgoCo saw limited resistance to the introduction of the Inventory Assortment Planning Tool, which enabled meaningful adoption of the tool. Many prior studies find considerable resistance to new algorithmic tools [e.g., 18, 50], so it is worth identifying the conditions under which we observed limited resistance. First, AlgoCo was founded after 2000 and had a data-first strategy and reputation. Individuals within the organization would have self-selected to work in an organization known for digital transformation. As such, buyers within AlgoCo might have been more willing to adopt new technologies than users in many other organizations. A second boundary condition is the approach taken by the data scientists at AlgoCo when developing the Inventory Planning Tool. These data scientists undertook a human-centered design process that focused on engagement, collaboration, and reskilling. This approach helped to ensure that the tool would be usable by buyers, again facilitating adoption. Relatedly, as part of this process, the data scientists ensured that buyers found the tool useful. They created features such as visualizations and automatic calculations of metrics that the buyers valued. The final boundary condition is occupational status of the end user group. Buyers within AlgoCo are a relatively high-status group, and prior research on digitization [e.g., 54, 73] shows that high status groups are more likely to undergo reskilling and adopt new technologies than low status groups who are more likely to undergo deskilling and in some cases replacement.

5.4 Implications for design

This study offers design implications for developers and data scientists: algorithms can be unnecessarily constrained by organizational structures, meaning developers might benefit from examining the organizational structures shaping the work of their users, including their role structures and organizational hierarchies. Practically, developers should consider how departmental targets - and thus decisions - are divided to accommodate human decision-making and what might be unnecessarily segmented. In this work, decisions were divided arbitrarily by clothing type, material, price point, and potential use. In other organizations, we have seen targets divided by geography, customer industry segment, and customer size, among others. Decisions segmented by organization chart may also go beyond targets to include things like technologies developed within an engineering organization. Such divisions may have implications both for how the decision space of an algorithm is circumscribed and how data is tagged and thus considered by the algorithm.

Additionally, our work highlights another reason for the value of user-centered design. In the case of AlgoCo, the data scientists co-designed the algorithm with the buyers, helping ensure buy-in and that the group handled the resulting tension with the organization chart collectively. Our study also reinforces the need for developers to consider and integrate the proper set of stakeholders; at the start of the design process, the algorithm was intended only for front-line buyers but it quickly implicated other roles and also buying managers and directors.

This work is also a reminder to leaders of organizations: organizations need to be designed with an eye toward algorithmic stakeholders as well. Changes to the organization chart may be needed to facilitate people leveraging the insights that algorithms may provide. Organization charts create clear separations of roles, but they also create incentives to maximize performance at one's own "leaf node" of the organization chart. Algorithms have no such constraints, but, without collaboration across the organization, the algorithm may be hamstrung.

Like the organization chart in this study, many ways of organizing departments have existed for decades, in some cases as long as such functions have existed. Such entrenched structures are likely difficult to change. Though not the focus of this work, and thus a decidedly incomplete perspective on engendering such change, AlgoCo's approach may offer some practical starting points. First, individual buyers and managers within the Merchandising Department were involved in discovering how the organization chart was constraining decisions. Specifically, data scientists designed the tool for use at the "leaf node" but allowed for enough flexibility to show the impact of optimizing at this level on aggregate performance metrics. Second, managers were allowed to "experiment" with a new decision-making process prior to ratifying a new organizational structure. Such trial periods may allow for mutual accommodation between the algorithm and the organization chart, ensuring that unforeseen issues with organization chart changes are addressed, and may allow individuals to become comfortable with the changes gradually.

6 CONCLUSION

Our research draws attention to an under-explored space in collaborative work and decision-making and highlights future opportunities to look beyond the impact of algorithms on individuals to the impact of algorithms on organizations. Our observations suggest that algorithms are likely to surface previously 'taken-for-granted' ways of organizing well beyond the scope of any one individual's work.

Though researchers have theorized that algorithm use should have implications for organizational structures and role interactions [7, 76], there has been limited research documenting how algorithms are coming into conflict with organizational structures. In this study, we showed how a company's organization chart arbitrarily segmented decision spaces, constrained the exploration of alternatives in decision-making, prioritized consideration of only one dimension of decision-making, and created hierarchical relationships between variables that were not hierarchical. Our study broke new conceptual ground by showing how one organization came to recognize and address these tensions.

REFERENCES

[1] Howard E. Aldrich. 1972. Technology and Organizational Structure: A Reexamination of the Findings of the Aston Group. *Administrative Science Quarterly* 17, 1 (March 1972), 26–43. https://doi.org/10.2307/2392089

[2] Callen Anthony. 2021. When Knowledge Work and Analytical Technologies Collide: The Practices and Consequences of Black Boxing Algorithmic Technologies. *Administrative Science Quarterly* 66, 4 (Dec. 2021), 1173–1212. https://doi.org/10.1177/00018392211016755

[3] Diane E. Bailey and Stephen R. Barley. 2020. Beyond design and use: How scholars should study intelligent technologies. *Information and Organization* 30, 2 (2020), 100286. https://doi.org/10.1016/j.infoandorg.2019.100286

[4] Liam J. Bannon and Kjeld Schmidt. 1989. CSCW - Four Characters in Search of a Context. *DAIMI Report Series* 18, 289 (Sept. 1989). https://doi.org/10.7146/dpb.v18i289.6667

[5] Stephen R. Barley. 1986. Technology as an Occasion for Structuring: Evidence from Observations of CT Scanners and the Social Order of Radiology Departments. *Administrative Science Quarterly* 31, 1 (March 1986), 78. https://doi.org/10.2307/2392767

[6] Stephen R. Barley. 1990. The Alignment of Technology and Structure through Roles and Networks. *Administrative Science Quarterly* 35, 1 (1990), 61–103. https://doi.org/10.2307/2393551

[7] Stephen R. Barley. 2020. *Work and Technological Change*. OUP Oxford.

[8] Stephen R. Barley, William H. Dutton, Sara Kiesler, Paul Resnick, Robert E. Kraut, and JoAnne Yates. 2004. Does CSCW need organization theory?. In *Proceedings of the 2004 ACM conference on Computer supported cooperative work*. ACM, Chicago Illinois USA, 122–124. https://doi.org/10.1145/1031607.1031628

[9] Chester I. Barnard. 1968. *The Functions of the Executive* (30th anniversary ed.). Harvard University Press, Cambridge, Massachusetts. Originally published in 1938.

[10] Nicholas Bloom, Luis Garicano, Raffaella Sadun, and John Van Reenen. 2014. The Distinct Effects of Information Technology and Communication Technology on Firm Organization. *Management Science* 60, 12 (Dec. 2014), 2859–2885. https://doi.org/10.1287/mnsc.2014.2013

[11] Sarah Brayne and Angèle Christin. 2021. Technologies of Crime Prediction: The Reception of Algorithms in Policing and Criminal Courts. *Social Problems* 68, 3 (Aug. 2021), 608–624. https://doi.org/10.1093/socpro/spaa004

[12] Taina Bucher. 2018. *If...Then*. Vol. 1. Oxford University Press. https://doi.org/10.1093/oso/9780190493028.001.0001

[13] Zana Buçinca, Maja Barbara Malaya, and Krzysztof Z. Gajos. 2021. To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI in AI-assisted Decision-making. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (April 2021), 1–21. https://doi.org/10.1145/3449287

[14] Ángel Alexander Cabrera, Adam Perer, and Jason I. Hong. 2023. Improving Human-AI Collaboration With Descriptions of AI Behavior. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1 (April 2023), 1–21. https://doi.org/10.1145/3579612

[15] Carrie J. Cai, Samantha Winter, David Steiner, Lauren Wilcox, and Michael Terry. 2019. "Hello AI": Uncovering the Onboarding Needs of Medical Practitioners for Human-AI Collaborative Decision-Making. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 1–24. https://doi.org/10.1145/3359206

[16] Shiye Cao and Chien-Ming Huang. 2022. Understanding User Reliance on AI in Assisted Decision-Making. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (Nov. 2022), 1–23. https://doi.org/10.1145/3555572

[17] Lingwei Cheng and Alexandra Chouldechova. 2022. Heterogeneity in Algorithm-Assisted Decision-Making: A Case Study in Child Abuse Hotline Screening. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (Nov. 2022), 1–33. https://doi.org/10.1145/3555101

[18] Angele Christin. 2017. Algorithms in practice: Comparing web journalism and criminal justice. *Big Data & Society* 4, 2 (2017), 1–14. https://doi.org/10.1177/2053951717718855

[19] Angèle Christin. 2018. Counting Clicks: Quantification and Variation in Web Journalism in the United States and France. *Amer. J. Sociology* 123, 5 (2018), 1382–1415. https://doi.org/10.1086/696137

[20] Amy C. Edmondson and Stacy E. McManus. 2007. Methodological fit in management field research. *Academy of Management Review* 32, 4 (2007), 1246–1264. https://doi.org/10.5465/amr.2007.26586086

[21] Lucca Eloy, Cara Spencer, Emily Doherty, and Leanne Hirshfield. 2023. Capturing the Dynamics of Trust and Team Processes in Human-Human-Agent Teams via Multidimensional Neural Recurrence Analyses. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1 (April 2023), 1–23. https://doi.org/10.1145/3579598

[22] Virginia Eubanks. 2017. *Automating inequality: how high-tech tools profile, police, and punish the poor* (first edition ed.). St. Martin's Press, New York, NY.

[23] Samer Faraj, Stella Pachidi, and Karla Sayegh. 2018. Working and organizing in the age of the learning algorithm. *Information and Organization* 28, 1 (2018), 62–70. https://doi.org/10.1016/j.infoandorg.2018.02.005

[24] Neil Fligstein. 1990. *The Transformation of Corporate Control*. Harvard University Press, Cambridge, MA.

[25] Asbjørn Ammitzbøll Flügge, Thomas Hildebrandt, and Naja Holten Møller. 2021. Street-Level Algorithms and AI in Bureaucratic Decision-Making: A Caseworker Perspective. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (April 2021), 1–23. https://doi.org/10.1145/3449114

[26] Asbjørn William Ammitzbøll Flügge, Thomas Hildebrandt, and Naja Holten Møller. 2020. Algorithmic Decision Making in Public Services: A CSCW-Perspective. In *Companion of the 2020 ACM International Conference on Supporting Group Work*. ACM, Sanibel Island Florida USA, 111–114. https://doi.org/10.1145/3323994.3369886

[27] Riccardo Fogliato, Alexandra Chouldechova, and Zachary Lipton. 2021. The Impact of Algorithmic Risk Assessments on Human Predictions and its Analysis via Crowdsourcing Studies. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (Oct. 2021), 1–24. https://doi.org/10.1145/3479572

[28] Jay R. Galbraith. 1974. Organization Design: An Information Processing View. *Interfaces* 4, 3 (1974), 28–36. http://www.jstor.org/stable/25059090

[29] Tarleton Gillespie. 2014. The Relevance of Algorithms. In *Media Technologies*, Tarleton Gillespie, Pablo J. Boczkowski, and Kirsten A. Foot (Eds.). The MIT Press, 167–194. https://doi.org/10.7551/mitpress/9780262525374.003.0009

[30] Nina Grgić-Hlača, Christoph Engel, and Krishna P. Gummadi. 2019. Human Decision Making with Machine Assistance: An Experiment on Bailing and Jailing. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 1–25. https://doi.org/10.1145/3359280

[31] Luther Gulick and Lyndall Urwick. 2003. *Papers on the Science of Administration*. Early Sociology of Management and Organizations, Vol. 2nd ed. Routledge. Originally published in 1937.

[32] Rebecca M. Henderson and Kim B. Clark. 1990. Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms. *Administrative Science Quarterly* 35, 1 (March 1990), 9–30. https://doi.org/10.2307/2393549

[33] Pamela Hinds and Sara Kiesler. 1995. Communication across Boundaries: Work, Structure, and Use of Communication Technologies in a Large Organization. *Organization Science* 6, 4 (Aug. 1995), 373–393. https://doi.org/10.1287/orsc.6.4.373

[34] Kenneth Holstein, Maria De-Arteaga, Lakshmi Tumati, and Yanghuidi Cheng. 2023. Toward Supporting Perceptual Complementarity in Human-AI Collaboration via Reflection on Unobservables. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1 (April 2023), 1–20. https://doi.org/10.1145/3579628

[35] Yoyo Tsung-Yu Hou and Malte F. Jung. 2021. Who is the Expert? Reconciling Algorithm Aversion and Algorithm Appreciation in AI-Supported Decision Making. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (Oct. 2021), 1–25. https://doi.org/10.1145/3479864

[36] Christopher Hutchison. 1994. Patterns of language in organizations: implications for CSCW. In *Design Issues in CSCW*. Springer, 89–117.

[37] Katherine C. Kellogg, Melissa A. Valentine, and Angèle Christin. 2020. Algorithms at Work: The New Contested Terrain of Control. *Academy of Management Annals* 14, 1 (2020), 366–410. https://doi.org/10.5465/annals.2018.0174

[38] Kalle Kusk and Midas Nouwens. 2022. Platform-Mediated Food Delivery Work: A Review for CSCW. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (Nov. 2022), 1–25. https://doi.org/10.1145/3555645

[39] Sarah Lebovitz, Hila Lifshitz-Assaf, and Natalia Levina. 2022. To Engage or Not to Engage with AI for Critical Judgments: How Professionals Deal with Opacity When Using AI for Medical Diagnosis. *Organization Science* 33, 1 (Jan. 2022), 126–148. https://doi.org/10.1287/orsc.2021.1549

[40] Min Kyung Lee, Daniel Kusbit, Evan Metsky, and Laura Dabbish. 2015. Working with Machines: The Impact of Algorithmic and Data-Driven Management on Human Workers. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, Seoul Republic of Korea, 1603–1612. https://doi.org/10.1145/2702123.2702548

[41] Caitlin Lustig and Bonnie Nardi. 2015. Algorithmic Authority: The Case of Bitcoin. In *2015 48th Hawaii International Conference on System Sciences*. IEEE, HI, USA, 743–752. https://doi.org/10.1109/HICSS.2015.95

[42] Caitlin Lustig, Katie Pine, Bonnie Nardi, Lilly Irani, Min Kyung Lee, Dawn Nafus, and Christian Sandvig. 2016. Algorithmic Authority: the Ethics, Politics, and Economics of Algorithms that Interpret, Decide, and Manage. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, San Jose California USA, 1057–1062. https://doi.org/10.1145/2851581.2886426

[43] James G. March and Herbert A. Simon. 1958. *Organizations*. John Wiley & Sons, Inc, New York.

[44] Melissa Mazmanian and Christine M. Beckman. 2018. "Making" Your Numbers: Engendering Organizational Control Through a Ritual of Quantification. *Organization Science* 29, 3 (June 2018), 357–379. https://doi.org/10.1287/orsc.2017.1185

[45] Henry Mintzberg. 1973. *The nature of managerial work*. Harper & Row, New York, NY.

[46] Sauro Mocetti, Marcello Pagnini, and Enrico Sette. 2017. Information Technology and Banking Organization. *Journal of Financial Services Research* 51, 3 (June 2017), 313–338. https://doi.org/10.1007/s10693-016-0244-3

[47] Lawrence B. Mohr. 1971. Organizational Technology and Organizational Structure. *Administrative Science Quarterly* 16, 4 (Dec. 1971), 444–459. https://doi.org/10.2307/2391764

[48] Katelyn Morrison, Donghoon Shin, Kenneth Holstein, and Adam Perer. 2023. Evaluating the Impact of Human Explanation Strategies on Human-AI Visual Decision-Making. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1 (April 2023), 1–37. https://doi.org/10.1145/3579481

[49] Imani Munyaka, Zahra Ashktorab, Casey Dugan, J. Johnson, and Qian Pan. 2023. Decision Making Strategies and Team Efficacy in Human-AI Teams. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1 (April 2023), 1–24. https://doi.org/10.1145/3579476

[50] Stella Pachidi, Hans Berends, Samer Faraj, and Marleen Huysman. 2021. Make Way for the Algorithms: Symbolic Actions and Change in a Regime of Knowing. *Organization Science* 32, 1 (Jan. 2021), 18–41. https://doi.org/10.1287/orsc.2020.1377

[51] Hyanghee Park, Daehwan Ahn, Kartik Hosanagar, and Joonhwan Lee. 2022. Designing Fair AI in Human Resource Management: Understanding Tensions Surrounding Algorithmic Evaluation and Envisioning Stakeholder-Centered Solutions. In *CHI Conference on Human Factors in Computing Systems*. ACM, New Orleans LA USA, 1–22. https://doi.org/10.1145/3491102.3517672

[52] Joon Sung Park, Karrie Karahalios, Niloufar Salehi, and Motahhare Eslami. 2022. Power Dynamics and Value Conflicts in Designing and Maintaining Socio-Technical Algorithmic Processes. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW1 (March 2022), 1–21. https://doi.org/10.1145/3512957

[53] Sharon K. Parker and Gudela Grote. 2022. Automation, Algorithms, and Beyond: Why Work Design Matters More Than Ever in a Digital World. *Applied Psychology* 71, 4 (Oct. 2022), 1171–1204. https://doi.org/10.1111/apps.12241

[54] Roger Penn, Michael Rose, and Jill Rubery (Eds.). 1994. *Skill and occupational change.* Oxford University Press, Oxford ; New York.

[55] Anette C. M. Petersen, Lars Rune Christensen, Richard Harper, and Thomas Hildebrandt. 2021. "We Would Never Write That Down": Classifications of Unemployed and Data Challenges for AI. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (April 2021), 1–26. https://doi.org/10.1145/3449176

[56] Jeffrey Pfeffer and Huseyin Leblebici. 1977. Information Technology and Organizational Structure. *Pacific Sociological Review* 20, 2 (April 1977), 241 – 261.

[57] Bogdana Rakova, Jingying Yang, Henriette Cramer, and Rumman Chowdhury. 2021. Where Responsible AI meets Reality: Practitioner Perspectives on Enablers for Shifting Organizational Practices. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (April 2021), 1–23. https://doi.org/10.1145/3449081

[58] Charvi Rastogi, Yunfeng Zhang, Dennis Wei, Kush R. Varshney, Amit Dhurandhar, and Richard Tomsett. 2022. Deciding Fast and Slow: The Role of Cognitive Biases in AI-assisted Decision-making. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW1 (March 2022), 1–22. https://doi.org/10.1145/3512930

[59] Daniela Retelny, Sébastien Robaszkiewicz, Alexandra To, Walter S. Lasecki, Jay Patel, Negar Rahmati, Tulsee Doshi, Melissa Valentine, and Michael S. Bernstein. 2014. Expert crowdsourcing with flash teams. In *Proceedings of the 27th annual ACM symposium on User interface software and technology.* ACM, Honolulu Hawaii USA, 75–85. https://doi.org/10.1145/2642918.2647409

[60] Samantha Robertson, Tonya Nguyen, and Niloufar Salehi. 2021. Modeling Assumptions Clash with the Real World: Transparency, Equity, and Community Challenges for Student Assignment Algorithms. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems.* ACM, Yokohama Japan, 1–14. https://doi.org/10.1145/3411764.3445748

[61] Devansh Saxena, Karla Badillo-Urquiola, Pamela J. Wisniewski, and Shion Guha. 2021. A Framework of High-Stakes Algorithmic Decision-Making for the Public Sector Developed through a Case Study of Child-Welfare. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (Oct. 2021), 1–41. https://doi.org/10.1145/3476089

[62] Beau G. Schelble, Christopher Flathmann, Nathan J. McNeese, Guo Freeman, and Rohit Mallick. 2022. Let's Think Together! Assessing Shared Mental Models, Performance, and Trust in Human-Agent Teams. *Proceedings of the ACM on Human-Computer Interaction* 6, GROUP (Jan. 2022), 1–29. https://doi.org/10.1145/3492832

[63] Beau G. Schelble, Christopher Flathmann, Geoff Musick, Nathan J. McNeese, and Guo Freeman. 2022. I See You: Examining the Role of Spatial Information in Human-Agent Teams. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (Nov. 2022), 1–27. https://doi.org/10.1145/3555099

[64] Kjeld Schmidt and Liam Bannon. 1992. Taking CSCW seriously: Supporting articulation work. *Computer Supported Cooperative Work (CSCW)* 1, 1-2 (March 1992), 7–40. https://doi.org/10.1007/BF00752449

[65] Herbert A. Simon. 1957. *Models of man, social and rational. Mathematical essays on rational human behavior in a social setting.* J. Wiley & Sons.

[66] Herbert A. Simon. 1991. Bounded Rationality and Organizational Learning. *Organization Science* 2, 1 (1991), 125–134. http://www.jstor.org/stable/2634943

[67] Herbert A. Simon. 1997. *Administrative Behavior* (fourth ed.). Free Press, New York, NY. Original work published 1945.

[68] Herbert A. Simon. 2019. *The Sciences of the Artificial* (reissue of the third edition with a new introduction ed.). The MIT Press. https://direct.mit.edu/books/book/4551/The-Sciences-of-the-Artificial Original work published in 1968.

[69] Harini Suresh and John Guttag. 2021. A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle. In *Equity and Access in Algorithms, Mechanisms, and Optimization.* ACM, – NY USA, 1–9. https://doi.org/10.1145/3465416.3483305

[70] James D. Thompson. 2003. *Organizations in action: Social science bases of administrative theory.* Transaction Publishers, New Brunswick, New Jersey. Original work published in 1967.

[71] Melissa A. Valentine and Rebecca Hinds. 2022. How Algorithms Change Occupational Expertise by Prompting Explicit Articulation and Testing of Experts' Theories. https://doi.org/10.2139/ssrn.4246167

[72] Melissa A. Valentine, Daniela Retelny, Alexandra To, Negar Rahmati, Tulsee Doshi, and Michael S. Bernstein. 2017. Flash Organizations: Crowdsourcing Complex Work by Structuring Crowds As Organizations. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems.* ACM, Denver Colorado USA, 3523–3537. https://doi.org/10.1145/3025453.3025811

[73] Steven Peter Vallas. 1990. The Concept of Skill: A Critical Review. *Work and Occupations* 17, 4 (Nov. 1990), 379–398. https://doi.org/10.1177/0730888490017004001

[74] Elmira Van Den Broek, Anastasia Sergeeva, and Marleen Huysman Vrije. 2021. When the Machine Meets the Expert: An Ethnography of Developing AI for Hiring. *MIS Quarterly* 45, 3 (Sept. 2021), 1557–1580. https://doi.org/10.25300/

MISQ/2021/16559

[75] Helena Vasconcelos, Matthew Jörke, Madeleine Grunde-McLaughlin, Tobias Gerstenberg, Michael S. Bernstein, and Ranjay Krishna. 2023. Explanations Can Reduce Overreliance on AI Systems During Decision-Making. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1 (April 2023), 1–38. https://doi.org/10.1145/3579605

[76] Georg Von Krogh. 2018. Artificial Intelligence in Organizations: New Opportunities for Phenomenon-Based Theorizing. *Academy of Management Discoveries* 4, 4 (Dec. 2018), 404–409. https://doi.org/10.5465/amd.2018.0084

[77] Lauren Waardenburg, Marleen Huysman, and Anastasia V. Sergeeva. 2022. In the Land of the Blind, the One-Eyed Man Is King: Knowledge Brokerage in the Age of Learning Algorithms. *Organization Science* 33, 1 (Jan. 2022), 59–82. https://doi.org/10.1287/orsc.2021.1544

[78] Lauren Waardenburg, Anastasia Sergeeva, and Marleen Huysman. 2022. Juggling Street Work and Data Work: An Ethnography of Policing and Reporting Practices. *Academy of Management Proceedings* 2022, 1 (Aug. 2022), 16697. https://doi.org/10.5465/AMBPP.2022.215

[79] Oliver E. Williamson. 1970. *Corporate Control and Business Behavior: An inquiry into the effects of organizational form on enterprise behavior.* Prentice-Hall, Inc., Englewood Cliffs, New Jersey.

[80] Christine Wolf and Jeanette Blomberg. 2019. Evaluating the Promise of Human-Algorithm Collaborations in Everyday Work Practices. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 1–23. https://doi.org/10.1145/3359245

[81] JoAnne Yates. 1993. *Control through communication: the rise of system in American management* (johns hopkins paperbacks ed.). Johns Hopkins Univ. Press, Baltimore, Md. Originally published in 1989.

Received 18 July 2023; Revised 30 April 2024