

Management Style Under the Spotlight: Evidence from Studio Recordings*

Girum Abebe [†]
Marcel Fafchamps [‡]
Michael Koelle [§]
Simon Quinn [¶]
Tom Schwantje ^{||}

January 16, 2026

Abstract

We introduce a new method for measuring managerial traits of young professionals: using management vignettes in a video studio. This method – analysed through the lens of a Bayesian hierarchical model – allows us to identify four distinct managerial archetypes (which we term ‘rule-based’, ‘affiliative’, ‘power-based’ and ‘discretion-based’). We find that past labour market exposure (including exposure induced through a previous field experiment) correlates strongly with the propensity to act as a rule-based manager. We then use the videos to run an incentivised experiment with firm managers, to elicit preferences over young professionals. Strikingly, we find that firms consistently prefer the rule-based managerial style for entry-level managerial positions. Empirically, our results highlight an underexplored mechanism for labour market exclusion among young professionals. Methodologically, we demonstrate the value of controlled vignette scenarios for assessing managerial traits. Our findings underscore the importance of managerial training in shaping labour market outcomes, and offer new avenues for studying the development of managerial talent.

*We are grateful to Jishnu Das, David Deming, Pascaline Dupas, Stephen Hansen, Seema Jayachandran, Nicola Limodio, Paula López-Peña, Fiker Negash, Pia Pinger, Andrea Prat, Imran Rasul, Bruce Sacerdote, Rafaella Sadun, John van Reenen, Duncan Webb and Chris Woodruff for helpful comments. We thank seminar audiences at the CEPR Paris Symposium, City St George’s (University of London), CSEF-IGIER Symposium on Economics and Institutions, the NBER Organizational Economics Working Group, Princeton University, Stanford University and Queen’s University (Firms and Inclusive Growth Workshop). This project was funded by the Private Enterprise Development in Low-Income Countries programme (PEDL). We also want to thank Gezahegn Gbremedhin for outstanding research assistance. The project would not have been possible without the constant support of Rose Page and the Centre for the Study of African Economies (University of Oxford), nor without the support of EconInsights in Addis Ababa. IRB approval was obtained from the University of Oxford (ECONCIA21-22-19). We would like to acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work: <http://dx.doi.org/10.5281/zenodo.22558>.

[†]**International Finance Corporation:** gtefera@ifc.org

[‡]**Stanford University:** fafchamp@stanford.edu.

[§]**OECD:** michael.koelle@oecd.org.

[¶]**Imperial College London:** simon.quinn@imperial.ac.uk.

^{||}**Bocconi University:** tom.schwantje@unibocconi.it.

1 Introduction

The late twentieth century witnessed a marked expansion of formal education in many low- and middle-income countries (Alesina, Hohmann, Michalopoulos, and Papaioannou, 2021; Asher, Novosad, and Rafkin, 2022). As a result, urban labour markets in low- and middle-income countries have seen a marked change in the educational composition of new entrants, with a substantial share of young job seekers having significantly higher levels of formal education than their parents. This educational mobility has important implications for understanding barriers to labour market inclusion – particularly among those aspiring to professional jobs that, for many, differ markedly from the roles held by their parents. At the same time, youth in developing countries struggle to access high-quality jobs after completing their higher education (Abebe, Caria, and Ortiz-Ospina, 2021a; Carranza, Garlick, Orkin, and Rankin, 2022). While the role of formal schooling in the accumulation of human capital, including the development of non-cognitive skills, is well documented (Bowles and Gintis, 1977, 2002; Heckman and Rubinstein, 2001), little is known about what attributes employers demand from educated young professionals – and how prior labour market experience and family background shape those attributes. Moreover, measurement of non-cognitive work skills is challenging, and such skills are often only observed – if at all – for those individuals who are actually employed.

To make progress on this issue, we design and implement a novel style of controlled experiment to measure non-cognitive professional attributes at the individual level. Specifically, we focus on managerial traits, which are important ingredients to the allocation of talent in modernising economies (Bandiera, Hansen, Prat, and Sadun, 2020; Dahlstrand, László, Schweiger, Bandiera, Prat, and Sadun, 2025). We use this measure to test whether such professional traits constitute a barrier to labour market inclusion for educated youth, as highlighted by Abebe, Fafchamps, Koelle, and Quinn (2024) in the same context. Specifically, we work with a large sample of young professionals, chosen for their interest in business and entrepreneurship – many of whom are not yet working in managerial roles. We invite these young professionals to a studio, where we explain that they will each watch a series of actors portraying different managerial scenarios. We video-record each participant’s response to each vignette – explaining that, with their permission, their video will be assessed by human resource managers for their quality as an entry-level manager and as an entrepreneur. We run this assessment using HR managers in established firms across a range of industries – using an incentive-compatible mechanism to elicit the firms’ assessments. Leveraging this controlled experimental setting, we are able to pose the same set of managerial challenges to every respondent in this sample – avoiding the endogenous assignment of managerial tasks to managers that is otherwise inherent in working with observational data on managerial decisions.

Central to our approach is the notion that managerial styles are often best understood as complex bundles of behaviours, rather than being straightforward to observe and record (Goleman, 2000; Benson and Shaw, 2025). For this reason, we deliberately chose a diverse range of managerial problems (rather

than, for example, repeating a single class of scenarios), and we encouraged open-ended answers (rather than, for example, forcing our respondents to choose from a small menu of available actions). Similarly, we then encode respondents' answers in a high-dimensional space (specifically, we measure responses through the combination of managerial action, justification, tone, and source of authority). To analyse this data, we build and estimate a bespoke Bayesian hierarchical model. This has the flavour of Latent Dirichlet Allocation ('LDA') models (Bandiera et al., 2020; Blei, Ng, and Jordan, 2003; Griffiths, 2004) and – like LDA models – allows us to characterise heterogeneity in managerial styles across a high-dimensional response space.

We have three key results. First, our method – that is, the combination of a studio exercise with a Bayesian hierarchical model – succeeds in identifying meaningful and substantial differences in managerial traits across young professionals. Specifically, we describe four latent 'pure types' of managerial styles. After estimation, we label these management styles as '*rule-based*', '*affiliative*', '*power-based*' and '*discretion-based*'. These types differ quite radically in their conceptualisation of the role of a junior manager. Managers using a rule-based style see themselves as implementing directives set by more senior managers: they are much more likely than other types to rely upon formal policy, and they emphasise the firm's interests in doing so. Affiliative managers, on the other hand, seek shared ground – often yielding to their counterparty, and emphasising shared interests as they do so. Managers using a power-based style seem to view their role in terms of implementing their personal objectives, and do so relying upon their personal authority. Managers using a discretion-based style are less assertive and often rely on ad-hoc approaches to justify their decisions.

Second, we find a clear preference across firms for rule-based managers over the other three types. We designed our experiment with a view that different firms may have different preferences over different types of manager (see, in particular, Bandiera et al. (2020)). Our findings do not fully conform with this expectation: while we observe some heterogeneity in preferences for management styles across employers, firms seem largely united in their view that the rule-based style of management is best – both for entry-level managerial roles and for self-employment. After obtaining our main Bayesian estimates, we returned to the field to run semi-structured interviews with firm employees; these interviews show that employees share their managers' preferences for the rule-based style.

Third, young professionals' managerial behaviour is strongly related to their prior labour market experience. In particular, we find that rule-based type respondents are more likely than the other types to have been self-employed, to earn higher wages (conditional on wage employment), and to report higher reservation wages and reservation profits. Strikingly, respondents are significantly more likely to exhibit a rule-based management style if they were randomly treated in a previous field experiment that increased their exposure to management practices in medium and large firms (Abebe et al., 2024). This indicates that the relationship between labour market exposure and managerial style is causal. We find that this experimental effect is driven entirely by individuals whose parents did not finish primary school. This implies

that the impact of early-career labour market exposure is particularly valuable for those who are likely to be systematically excluded from labour market opportunities – and that lack of ‘professional socialisation’ is likely to be a key mechanism by which labour market inequality can persist across generations.

Together, these results contribute to three bodies of literature. First, our results complement the recent literature studying labour market exclusion among young professionals. Our results show that individuals with greater labour market exposure are more likely to exhibit the rule-based style of management – and, conversely, that the rule-based style of management is in higher demand among prospective employers. Together, these two results imply that embodied managerial traits may act as a key mechanism by which some young professionals enjoy sustained labour market success while others face exclusion. This supports a key insight from our earlier work (Abebe et al., 2024) – and, indeed, our results here help to interpret the mechanisms identified in that earlier paper. More generally, our results highlight the importance of non-cognitive skills for employment in urban low-income settings (Bassi and Nansamba, 2022); indeed, our result on the relevance of parental education suggests that embodied managerial traits may act as a form of ‘cultural capital’ (Bourdieu, 1986), by which existing labour market inequalities are reinforced (Zimmerman, 2019; Falk, Kosse, and Pinger, 2020; Barrios-Fernández, Neilson, and Zimmerman, 2024; Karna, List, Simon, and Uchida, 2025; Shukla, 2025; Attanasio, De Paula, and Toppeta, 2025).

Second, our findings add to our understanding of managerial traits. The past two decades have witnessed a substantial expansion in economists’ interest in the role of management. This has been driven by a new empirical focus on the measurement of management practices, spearheaded by the World Management Survey (Scur, Sadun, Van Reenen, Lemos, and Bloom, 2021). A related literature has highlighted the critical role of managers themselves, using comparisons both across firms (Bertrand and Schoar, 2003; Bloom, Brynjolfsson, Foster, Jarmin, Patnaik, Saporta-Eksten, and Van Reenen, 2019) and within firms (Hoffman and Tadelis, 2021; Lazear, Shaw, and Stanton, 2015; Metcalfe, Sollaci, and Syverson, 2023; Sen, 2024), as well as the high cost of good managers in developing countries (Hjort, Malmberg, and Schoellman, 2022). Managerial *traits* embodied in *individual managers* have been the focus of a smaller literature in economics (see, in particular, Malmendier, Tate, and Yan (2011), Kaplan, Klebanov, and Sorensen (2012), Bandiera, Guiso, Prat, and Sadun (2015), Benmelech and Frydman (2015), Bandiera et al. (2020) and López-Peña, Mozumder, Rabbani, and Woodruff (2025)). Much of this literature has studied decisions taken by managers, in their managerial capacity; that is, samples of respondents who (i) have already been selected by firms for their managerial abilities, and (ii) whose managerial challenges are determined by the particular managerial contexts in which they find themselves (as pointed out by Weidmann, Vecchi, Said, Deming, and Bhalotra, 2024). In contrast to this literature, by measuring managerial traits through a controlled studio environment, we are able to shed new light on the role of such traits as a mechanism for labour market exclusion. That is, we are able to measure managerial traits even among those who are not employed as managers.

On the one hand, our result that managerial traits are related to past labour market exposure is

broadly consistent with earlier results showing the relevance of past experiences in shaping managerial style (Malmendier et al., 2011; Benmelech and Frydman, 2015). In contrast, our result on the relative homogeneity of firm preferences provides a counterpoint to the seminal earlier work of Bandiera et al. (2020) – who find heterogeneous firm-manager match quality among different kinds of employed senior managers. This presents a novel insight on firm preferences over management: while different firms may have different preferences over the managerial styles of their *senior* leaders, firms seem reasonably homogeneous in their preferences for *entry-level* managers among a broad pool of young professionals.

Third, methodologically, our results contribute to recent literature using innovative and open-ended data collection methods in economics. This includes – in particular – recent literature on the relevance of tone and expression in economic communications (Chang, Dai, Feng, Han, Shi, and Zhang, 2025; Handlan and Sheng, 2023; Gorodnichenko, Pham, and Talavera, 2023), the use of photographs to proxy labour market potential (Guenzel, Kogan, Niessner, and Shue, 2025) and pre-trial detention (Ludwig and Mullainathan, 2023), and on the use of machine learning techniques to understand CEO performance (Borgschulte, Guenzel, Liu, and Malmendier, 2021). More generally, our approach complements recent advances in the use of open-ended survey questions in economics (Stantcheva, 2021; Haaland, Roth, Stantcheva, and Wohlfart, 2024). Separately, our paper illustrates and validates the use of studio vignettes as a credible method for eliciting managerial traits in field settings. This builds on a small literature in economics that has used a different kind of team-based ‘management simulations’, to test the impacts of managerial training in the field (Macchiavello, Menzel, Rabbani, and Woodruff, 2020) and to assess managerial potential in the lab (Weidmann et al., 2024; Weidmann and Deming, 2021; Weidmann, Xu, and Deming, 2025). More generally, we note that the use of video responses to pre-recorded questions is increasingly common as part of ‘assessment centres’ for job recruitment exercises and university entrance processes – but has received relatively little empirical attention in economics.

The paper proceeds as follows. In section 2, we describe the experimental context and implementation. In section 3, we use a Bayesian hierarchical model to characterise heterogeneity in management traits. We go on in that section to describe the characteristics of those types, and then provide a sentence-embedding analysis of different types’ vocabulary. In section 4, we turn to the firm side – to characterize firm preferences over different managerial approaches (including a robustness check exploiting an alternative source of exogenous variation – namely, random variation in the gender of the actors viewed by studio participants. We go on to report a model-informed field survey of workers. Section 5 shows the impact of our earlier management placement upon respondents’ management styles and firms’ assessments of those respondents. Section 6 concludes.

2 The experiment

2.1 Vignettes in the studio

Our experiment is designed to measure management traits among young professionals across a series of realistic management scenarios. To do this, we used a studio setting, in which respondents participated in a role-play scenario. Specifically, we ran a series of separate vignettes; in each vignette, the respondent watched a video of a paid actor, who played the role of a counter-party in a managerial problem. For each vignette, we played the video (explaining that it showed an actor), and then asked the respondent to provide a short response – as if she or he were in a managerial role, responding to the character in the video. We explained to the respondents that their video recordings would be played to human resource managers in successful Ethiopian firms. We recorded each vignette using both a male actor and a female actor (with different actors for the different vignettes); we randomly varied whether respondents viewed the male or female actor, and we exploit this random variation in section 4.2.

Specifically, we set each respondent five different scenarios:¹

- (i). *Line management of an employee*: The actor is an accounting clerk who has been absent for three days without prior notice. (S)he shows up for work on the fourth day, explaining that (s)he was unwell.
- (ii). *Negotiating with a supplier*: The actor plays a supplier, who explains that (s)he cannot fulfil an order according to agreed specifications, because of problems sourcing input materials. (S)he offers to supply a replacement of inferior quality instead. Respondents are told specifically that the firm they represent is known for producing and selling the highest quality products in the industry.
- (iii). *Negotiating a pay rise*: The actor is a production worker, who comes to ask for a pay rise. The worker argues that (i) within her unit, (s)he has been the most productive worker during the last three months, and one of the few people who exceeded the personal productivity targets; and (ii) (s)he has been with the company for ten years, and even her mother used to be an employee of the company. Respondents are told that the firm they represent has no plan to increase any salaries this year.
- (iv). *Negotiating an adjustment with the bank*: The actor is a bank manager, who calls to remind the firm about an unpaid loan installment (explaining that failure to pay increases interest payments and reduces the firm's credit rating). Respondents are told that the firm they represent will not have sufficient funds to pay the bank for another two weeks.

¹ The full scenario scripts are provided in Appendix A. The appendix is available online at http://www.simonrquinn.com/ManagementSpotlight_Appendix.pdf.

-
- (v). *Negotiating with a client*: The actor is a client, who has not paid what you have invoiced, in spite of one reminder. The client comes to place a new order.

2.2 Incentivised elicitation of firm preferences

In the second part of our experiment, we played the studio recordings to human resources managers from successful Ethiopian firms.² For each vignette, the human resources manager watched three responses, and was asked to rank these separately based on (i) the manager’s assessment of the respondents’ suitability as entry-level managers at their firm, and (ii) the manager’s assessment of the respondents’ suitability to run their own small business.

Specifically, each human resources manager was assigned to assess video recordings from three different respondents, for each of the five vignette scenarios. Each triplet of young professionals was assessed by at least two different managers for each vignette. To elicit revealed preferences comparing different candidates, we implemented two complementary direct elicitation mechanisms: one to elicit the perceived suitability of candidates for a managerial position in the firm; the other to elicit candidates’ perceived suitability to run their own business. We incentivised the assessment of suitability for managerial positions through the prospect of receiving respondents’ contact details (with respondents’ permission), and we incentivised the assessment of suitability to run a business through the prospect of the respondent being invited to a business plan competition. They were designed to be straightforward and to make truthful ranking an ‘obviously dominant’ strategy for respondents.³ Separately, we asked each human resources manager directly to assess the individual respondents (again, both for their suitability as an entry-level manager, and as an entrepreneur), using a series of stated-preference questions. We explain both the revealed-preference and stated-preference methods in detail in Appendix B.

2.3 Implementing our experiment

We conduct our experiment with 982 young professionals, who previously participated in an experiment in which a random subset was assigned to a management placement (see Abebe et al. (2024)).⁴ The respondents are primarily male (75%), highly educated (78% have a university degree), and relatively young (with an average age of 31). Most of the young professionals are in wage employment (68%) or self-employment (17%). Of those in wage employment, 98% work in a professional position, and 18% in a managerial position. Conditional on wage employment and self-employment, respectively, the median wage and profit are both 6,000 ETB (although self-employment income is more disperse across

² We ran this assessment with the senior member of the firm who was responsible for hiring. We refer to this person here as ‘the human resources manager’ – though this was not always the actual title or job description.

³ Both mechanisms closely resembled the ‘OSP-RSD’ ranking mechanism described by Li (2017).

⁴ We show robustness to attrition at the end of section 3.3.

respondents). This is approximately equal to the *average* household expenditure of 6,100 ETB in Addis Ababa.⁵

Despite our respondents' high level of education, there is significant variation in their socio-economic backgrounds. Only 11% of our respondents have a parent who completed a university degree; even more striking, 48% of respondents report that neither parent completed primary school.⁶ These patterns emphasise that these young professionals are participating in a labour market that differs markedly from the one their parents faced. Appendix Table A.32 summarises.

The sample of firms was drawn from a set of 713 established Ethiopian firms involved in a previous experiment (Abebe et al., 2024), supplemented by additional firms to account for attrition. These are medium-sized to large firms operating across a range of industries, in Addis Ababa and in the neighbouring towns of Bishoftu and Adama. Appendix Table A.34 describes the sample included in the project. We conducted interviews with a sample of 576 firms, where we interview a manager responsible for human resource decisions in the firm. These firms have a median of 58 employees (mean = 323) and employ a median of five managers (mean = 12). The interviewed managers work primarily in human resources (40%) and administration (34%). They have been in their current position for, on average, eight years. 80% have a university degree and 72% report having formal management training. Appendix Table A.33 summarises.

3 Management styles in the studio

Our first experimental objective is to characterise the main management styles among young Ethiopian professionals. To do this, we build a bespoke Bayesian Hierarchical Model; this allows us to identify distinct clusters of managerial behaviour. To do this, we first employed two enumerators to watch all of the videos, and to encode each response along the following four dimensions:⁷

- (i). *Action*: Did the respondent agree or disagree with the actor's request? (For example, in the third scenario, did the respondent agree to the employee's request for a pay rise?)
- (ii). *Authority*: What source of authority did the respondent rely upon? Specifically, did the respondent rely upon (a) their formal authority as a manager, (b) their seniority or personal authority, (c) higher principles, such as an appeal to the appropriateness of the action, or (d) did the respondent not rely upon any source of authority.

⁵ See the Ethiopian Socioeconomic Survey (Ethiopian Statistical Service and World Bank, 2023), Table 7.1.

⁶ Alesina et al. (2021) show that, for urban Ethiopia generally, the 'upward intergenerational mobility' – that is, the probability of completing primary school conditional upon neither parent having completed primary school – is about 55% (see their Appendix Table E.II). For a broader sample of 28 countries, Ouedraogo and Syrichas (2021) find that individuals born in 2000 and residing in urban areas exhibit, on average, upward intergenerational mobility of 78%.

⁷ Appendix C provides a detailed description.

-
- (iii). *Justification*: How did the respondent justify her or his action? Specifically, did the respondent (a) provide no justification, (b) emphasise the interests of the firm they represent, (c) emphasise the interest of the other party, (d) emphasise the shared interest of the firm and the other party, or (e) emphasise their personal interest?
- (iv). *Tone*: What tone did the respondent use? Specifically, was it (a) calm/assured, (b) assertive, or (c) aggressive?

3.1 Management traits among Ethiopian young professionals

We begin, in Figure 1, by describing the average behaviour across our five vignettes. Specifically, Figure 1 shows one row for each vignette; across that row, we characterise the average behaviour across each dimension of management (as discussed, these are action, authority, justification and tone). We disaggregate the data by the two enumerators. We draw two conclusions from the figure. First, we find substantial heterogeneity *across* different vignettes – and this is reflected across the bundle of different managerial dimensions. For example, when dealing with an employee seeking a pay rise (vignette 3), most respondents refuse the request; to do so, they rely heavily on either a formal policy or on personal authority, and their tone can be relatively assertive. In contrast, when dealing with their bank (vignette 4), respondents are much more likely to accede to the request, to place almost no reliance upon formal policy, and generally to use a less assertive tone. Similarly, when dealing with an employee seeking a pay rise (vignette 3), respondents frequently invoke their shared interest, but this rarely happens when dealing with a supplier failing to deliver a high-quality good (vignette 2).

Second, there is substantial heterogeneity *within* different vignettes, across respondents. This is evident from the disaggregation across the separate enumerators in Figure 1. Despite assessing the videos independently, the enumerators describe the responses in broadly similar terms – implying that the observed differences in action, authority, justification and tone are driven by heterogeneous approaches by different young professionals, rather than by differences in enumerators’ perceptions.

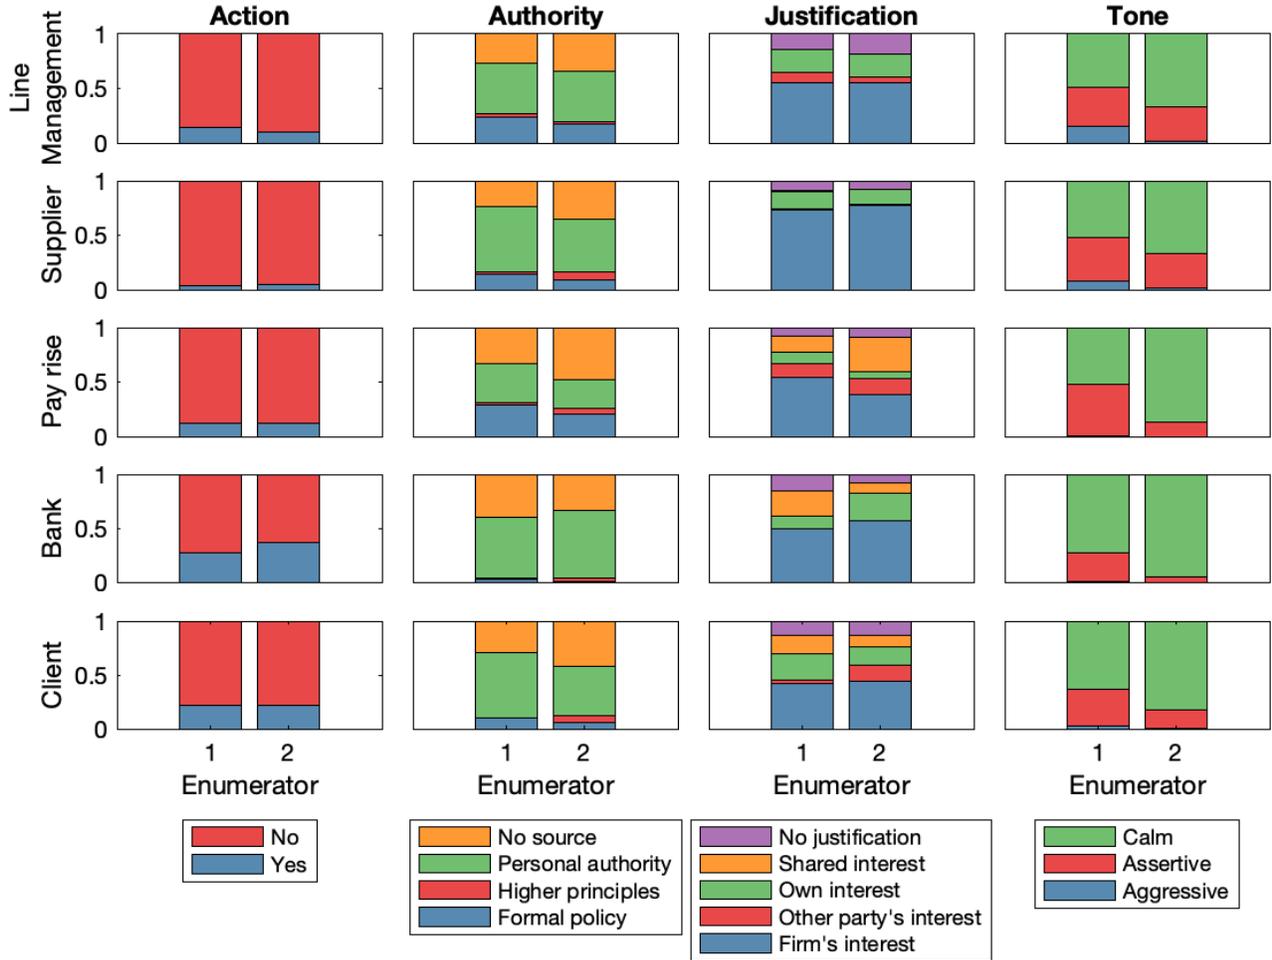
3.2 A Bayesian model of management styles

To extend this analysis, we now build a bespoke Bayesian Hierarchical Model. This model serves two related purposes. First, it allows us to cluster the different dimensions of managerial responses into different managerial styles. Second, the resulting classification will be a key input for understanding firm preferences in the following section.⁸

To do this, we specify a generative model in which each young professional – when responding to a given vignette – chooses a combination of action, authority, justification and tone. The model characterises

⁸ For this section, we only use the data from the enumerator encoding of the vignettes to estimate our model.

Figure 1: The distribution of behaviours by vignette



Notes This figure reports the distribution of the encoded behaviours by vignette and by enumerator. The bars show the probability that an enumerator encodes a specific behaviour for each dimension by vignette.

‘pure types’ that represent archetypal managerial styles; each young professional is then represented as a convex combination of those archetypes. A Bayesian Hierarchical framework is well suited for this task due to its ability to handle high-dimensional categorical data and to uncover latent themes. Very broadly, this is similar to several recent papers on heterogeneity in management styles (e.g. [Bandiera et al., 2020](#)).⁹

Specifically, we observe individuals $i \in \{1, \dots, N\}$ performing on vignettes $v \in \{1, \dots, 5\}$. For each individual assessment of a vignette, we have two enumerators, $e \in \{1, 2\}$. Each enumerator records a set of ‘attributes’ of the response (action, authority, justification and tone): $a \in \{1, \dots, 4\}$. Each attribute a has $J(a)$ possible categorical responses, $w_{ive}^a \in \{1, \dots, J(a)\}$.¹⁰ For each vignette, each respondent draws their behaviour from one of K pure types, z_{iv} . The total probability of the model is as follows:

$$P(W, Z, \theta, \phi, \zeta, \eta) = \underbrace{\prod_{j=1}^K P(\phi_j; \zeta)}_{\text{Dirichlet parameters}} \cdot \underbrace{\prod_{i=1}^N P(\theta_i; \eta)}_{\text{Types}} \cdot \prod_{v=1}^5 \sum_{k=1}^K \left(\underbrace{P(z_{iv} = k | \theta_i)}_{\text{Type assignment}} \cdot \underbrace{\prod_{e=1}^2 \prod_{a=1}^4 P(w_{ive}^a | \phi_{ka}, \psi_{av}, \chi_{ae})}_{\text{Studio behaviour}} \right), \quad (1)$$

where $P(\phi_j; \zeta)$ and $P(\theta_i; \eta)$ follow a Dirichlet distribution, $P(z_{iv} | \theta_i)$ follows a categorical distribution and $P(w_{ive}^a | \phi_{ka}, \psi_{av}, \chi_{ae})$ follows a Multinomial Logit distribution.

In this model, ϕ_{ka} characterises the behaviour of the pure types and θ_i characterises an individual as a convex combination of those archetypes. z_{iv} is an individual’s type for a specific vignette. ψ_{av} is a vignette fixed effect that allows average behaviour to vary across vignettes, and χ_{ae} is an enumerator fixed effect that allows average behaviour to vary across enumerators. We estimate using a Hamiltonian Monte Carlo algorithm in Stan ([Stan Development Team, 2024](#)); we discuss the parameterisation and estimation of this model in more detail in [Appendix E](#); standard MCMC diagnostics indicate excellent convergence.

3.3 Types of management

We estimate our model allowing for four ‘pure types’ of management. (In [Appendix F](#), we consider alternative versions with two, three and five types; we show that our general conclusions remain very similar.) [Figure 2](#) shows the four estimated ‘pure type’ management styles and their behaviour in our five vignettes:

- (i). Type 1 implements rules set by the firm. Specifically, they refuse the requests almost all of the time; in doing so, they are much more likely than the other types to rely upon formal policy – and much

⁹ Our model has some analogy to the class of Latent Dirichlet Allocation (‘LDA’) models. In the language of LDA models, we could think of each respondent-vignette pair as comprising a ‘document’, each document comprising four ‘words’, and each word being drawn respectively from one of four different dictionaries. Our inclusion of enumerator and vignette effects is a further variation on the typical LDA setup, akin to what [Biderman, Blei, Cai, Ciccarone, Feder, and Prat \(in progress\)](#) term a ‘fixed effects topic model’.

¹⁰ Specifically, $J(1) = 2$ for the action; $J(2) = 4$ for source of authority, $J(3) = 5$ for justification and $J(4) = 3$ for tone.

more likely to emphasise the firm's interests in doing so. They are more likely than the other three types to use an assertive tone; indeed, at times, they can even be aggressive.

- (ii). Type 2 seeks shared ground. Strikingly, the affiliative type is much more likely than the other three types to concede to their counterparty in the hypothetical discussion – notwithstanding that many of the vignette prompts suggested that such concessions were not in the interests of the hypothetical firm. This type tends to rely upon personal authority (or cites no specific authority), and is much more likely than the other three types to emphasise shared interests. Their tone is generally calm, but can be assertive.
- (iii). Type 3 implements their personal objectives. They refuse the requests almost all of the time; in doing so, they are overwhelmingly likely to rely on personal authority. As justification, they tend to emphasise the firm's interest – but are much more likely than the other three types to justify their actions by reference to their own personal interest. Similar to type 1, they tend to be assertive, or even aggressive.
- (iv). Type 4 does not provide a clear explanation for their decision. They refuse the requests almost all of the time – and, in doing so, are unlikely to rely on any source of authority, and often provide no specific justification for their decision. Of the four pure styles, this type of managers are most likely to be calm in their tone.

In Figure 3, we describe the distribution of individuals over types. To do this, we draw a tetrahedron, in which each vertex represents probability one of belonging to a particular type. We learn three things from this figure about our model results. First, the model distributes individuals quite evenly across the tetrahedron, such that any given individual exhibits a combination of different 'pure type' behaviours. Second, very few individuals sit close to the type 2 vertex: this implies that this style of behaviour, while empirically important, is something that individuals incorporate into their responses – rather than being a style that consistently describes any particular individual. Third, the stacked bar on the right of Figure 3 shows the observed vignette responses in terms of the average distribution of each pure type across individuals. We estimate that the most prevalent are type 1 (28.0%) and type 3 (27.9%); this is followed by type 4 (25.2%) and then type 2 (19.0%).

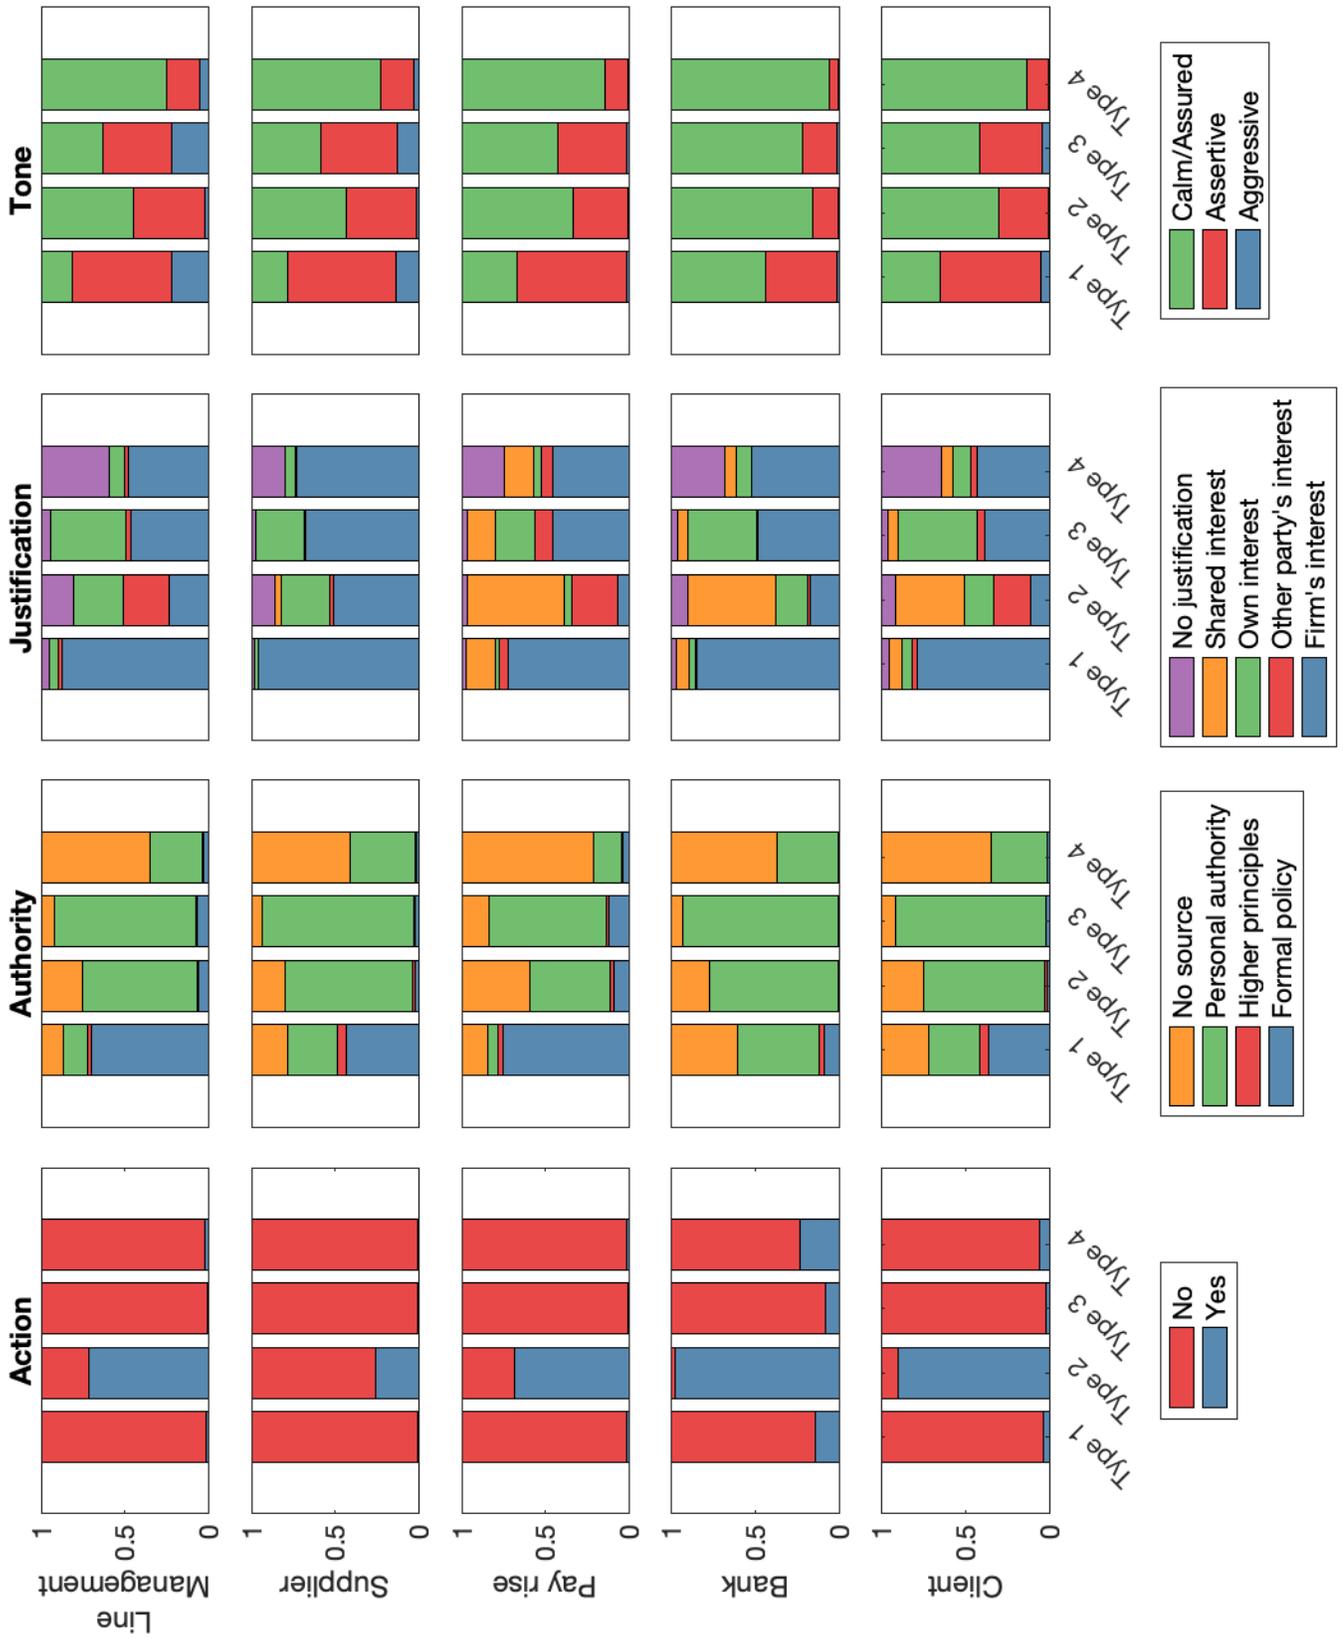
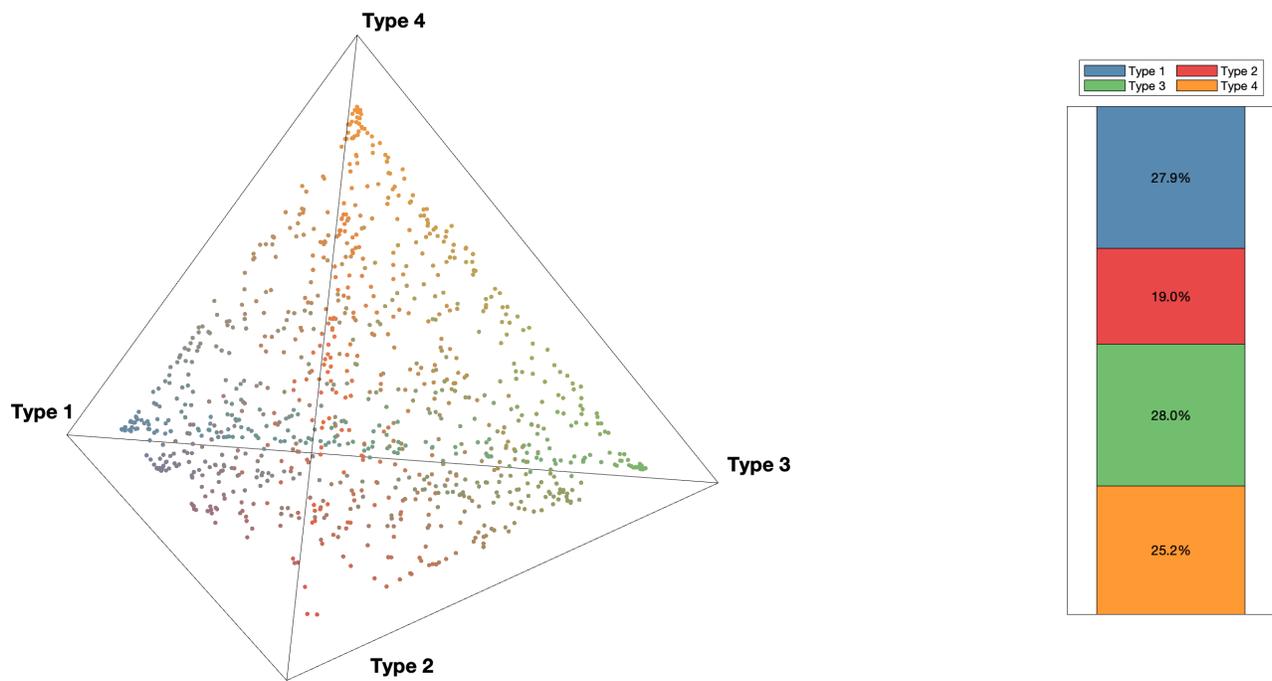


Figure 2: 'Pure type' management styles among Ethiopian young professionals

Figure 3: Distribution of types across individuals



About 60% of the original sample attended the studio. To compare the sample attending the studio with the original sample in the experiment, we run a number of descriptive logit estimations. In Appendix Table A.19, we first examine whether some individuals were more or less likely to attend the studio. First, in column 1, we show that – compared to the baseline sample – female participants are significantly less likely to attend, and that individuals with a BA degree are less likely to attend (although this is marginally significant). When we study attendance at the studio conditional on being in the four-year follow-up of the placement experiment (column 2), we similarly find that female participants are less likely to participate, and that individuals with an above-median reservation wage (conditional on employment) are less likely to attend. In Appendix Table A.20, we use an Inverse Probability Weighting to show the robustness of our main Bayesian estimates to concerns about selection on these covariates.¹¹

3.4 Managerial types and labour market experience

Having estimated the distribution of individuals across types, we now study the labour-market outcomes and trajectories of these individuals, conditional on their managerial traits. Panel A of Table 1 reports the average characteristics of individuals by type assignment. Type 1 is most likely to be male and self-employed and, relative to the other types, is more likely to report an above-median reservation wage and profit. The first type is also most likely to have completed at least a bachelor’s degree. Differences among the remaining three types are more subtle, yet several stand out. Type 2 is most likely to be female; and type 4 tends to report a below-median reservation wage.

Next, we examine how labour-market experience varies across managerial types. Approximately 12 months after respondents participated in the studio, we asked them about their labour-market experiences over the previous six years, including their employment status at six-month intervals. Panel B of Table 1 summarises these experiences by managerial type. The table indicates that type 1 managers spent more years in employment—particularly in permanent positions—and made about 0.2 fewer job transitions on average than other types, reflecting a more stable employment history. Finally, type 1 managers were also most likely to be in a management position when we interviewed them before participating in the studio.

¹¹ We find no evidence for the hypothesis that individuals in the treatment group of our original placement experiment are more likely to attend the studio. In column (3), we study whether any individual characteristics are predictive of treatment status conditional on attending the studio. We find suggestive evidence that female participants who are treated appear to be less likely to decline the studio invitation; the sample appears to be well-balanced on all other characteristics. Importantly for our subsequent analysis, we find good balance on treatment status and whether either parent finished primary school.

Table 1: Characteristics and Types: Summary Statistics

Panel A: Individual characteristics and types						
Type	Gender [1=male]	Wage employment indicator	Self- Employment indicator	Above Median Reservation Wage	Above Median Reservation Profit	At least a bachelor's degree
Type 1	.845	.709	.201	.449	.426	.838
Type 2	.771	.676	.15	.336	.305	.74
Type 3	.777	.681	.144	.334	.373	.765
Type 4	.797	.665	.147	.265	.298	.759
<i>p</i> -value	.047	.742	.215	<.001	<.001	.015

Panel B: Labour market experience and types					
Type	Employment <i>Years</i>	Permanent employment <i>Years</i>	Unemployed <i>Years</i>	Number of transitions <i>Count</i>	Management Position <i>Share</i>
Type 1	5.511	3.912	.335	.642	.186
Type 2	5.302	3.7	.455	.827	.124
Type 3	5.226	3.629	.513	.923	.141
Type 4	5.222	3.576	.576	.857	.11
<i>p</i> -value	.242	.068	.051	.001	.015

Notes This table describes the average characteristics of individuals (Panel A) and labour market experience (Panel B) of each type. Specifically, in Panel A this includes their gender (1 indicating male, 0 female), a dummy for their wage- and self-employment status, the probability they have an above-median reservation wage and profit based on data collected before respondents' attended the studio, and a dummy for whether or not they have included at least a bachelor's degree (BA, BSc, MA, MSc or PhD). Note that the median splits on reservation wage and profit do not yield a 50/50 split due to bunching in the underlying data at ETB10.000. Panel B includes the number of years they have been in employment including both self- and wage-employment, the number of years they have been in permanent employment, the number of years they have been unemployed, the number of labour market transitions they have gone through, and finally whether they were in a management position immediately before participating in the studio experiment. These numbers are calculated by assigning each individual to the pure type for which they have the highest estimated $\hat{\theta}_i$. Then, a conditional average is taken for each pure type. To test whether characteristics and labour market trajectories differ across types, we compute the Mahalanobis (Wald) statistic and compare it to a χ^2_3 distribution. This relies on the multivariate normal approximation to the posterior described in Section 4.1 of Gelman, Carlin, Stern, Dunson, Vehtari, and Rubin (2013). For non-binary outcomes, we first create a binary split at the median to split the sample.

3.5 Labelling the types

For ease of reference, we now add labels to the management styles exhibited by these four types. We stress that these labels reflect our suggested interpretation of the algorithmic output, and were not provided in advance to our unsupervised algorithm (as Ludwig and Mullainathan (2023, p.756) put it, ‘The algorithm discovers, and people name that discovery’). To assist in this process – and to provide further description of our types – we use natural language processing (‘NLP’). Specifically, we create Amharic transcriptions – and then English translations – using specialized Amharic NLP tools in collaboration with a private company (Hasab AI).¹² We use the English translations for text analysis, and then predict type membership using tokenized text (each token representing a combination of one or two words). We do this after first removing stopwords and common terms (but keeping negations). We train a multinomial logit with a ridge penalty (ℓ_2) in order to identify which tokens most strongly predict membership of each of our four types.¹³ We provide further details on this exercise in Appendix M; in that appendix, we draw similar conclusions from looking at the frequency with which respondents use individual words and type membership, and from text embedding analysis predicting type membership based on semantic word clusters (following, for example, Hansen, Sadun, Ramdas, and Fuller (2021) and Webb (2024)).

Figure 4 shows the results of our NLP analysis. Drawing on this figure and the results illustrated earlier in Figure 2, we label each type as follows:

- (i). Type 1: This type uses a *rule-based management style*: they refuse most requests, rely heavily on formal policy, and emphasise the firm’s interest (Figure 2). The NLP analysis (Figure 4, panel (a)) shows a clear emphasis on the importance of the firm as an organisation (‘organisation’, ‘company’), with a clear role for formal policy (‘regulations’, ‘procedures’, ‘policy’, ‘agreement’).
- (ii). Type 2: This type uses an *affiliative* management style: they are much more likely than any other type to accede to their counter-party’s request, and they often emphasise shared interests in doing so (Figure 2). The NLP analysis (Figure 4, panel (b)) reveals an emphasis on collaboration and the relationship (‘future’, ‘accepted’, ‘health’, ‘no problem’).
- (iii). Type 3: This type uses a *power-based* management style: respondents of this style are distinguished by their emphasis on personal authority (Figure 2); they seem to view executive power as flowing not so much from the rules of the firm as from the fact that they have been endowed with personal power. The NLP emphasises personal refusal (‘cannot’, ‘don’t accept’, ‘don’t’, ‘don’t want’), with a notable absence of tokens describing formal process (Figure 4, panel (c)).

¹² <https://www.hasab.ai/>

¹³ For this exercise, we treat the estimated latent-type vectors $\hat{\theta}_i$ as data; we do not propagate the estimation uncertainty in $\hat{\theta}_i$ into the text model.

(iv). Type 4: This type uses a *discretion-based* management style: these respondents provide limited justification for their decisions, and often provide no source of authority in support (Figure 2). For this type, the NLP highlights a much more diverse range of themes (‘wait’, ‘period’, ‘income’, ‘budget’: see panel (d) of Figure 4).

To shed further light on this characterisation, we also use an ‘AI enumerator’. Specifically, we upload each individual transcription to OpenAI’s server and pose a series of questions to GPT-5; these questions cover various attributes of each transcript, as well as asking for three adjectives to describe each response. This exercise reveals that the rule-based and affiliative management style are additionally associated with empathy – although the rule-based style is mainly likely to show empathy in vignettes involving external parties (i.e. the supplier, bank manager, client). The rule-based type also often mentions the future relationship, implementing a strict policy and using formal language in Amharic. The affiliative type is most likely to show respect, whereas the power-based type is most likely to express distrust. Appendix Table N shows the full results from this exercise and a breakdown by vignette, including the adjectives chosen by the AI enumerator to describe responses (which align closely with our earlier discussion). Finally, the AI enumerator reports that respondents were faithful to the hypothetical framing of the exercise, and did not (for example) try to impress by mentioning items relevant to their (real-world) CV; such references appear in slightly fewer than 0.2% of responses.

Finally, we can describe types in terms of respondents’ stated intentions. Prior to recording each video, we asked each respondent how they intended to answer, and why. Table 2 describes the average prevalence of different answers to these questions, across the four identified types.¹⁴ Specifically, we report prevalence of (i) distrust as a stated motivation for the planned response, (ii) relationship maintenance as a stated motivation for the planned response, (iii) intention to explain the decision in terms of formal company procedures, (iv) intention to set an example for future cases. These reported intentions align with our earlier interpretation of the four types. The affiliative type, for example, is least likely to distrust their counterparty, and much more likely to emphasise the value of maintaining an ongoing relationship; in doing so, they are far less likely than the other types to emphasise the importance of formal procedures. In contrast, the rule-based managerial style stands out for its relatively high distrust of the counterparty and – above all – for being much more likely than the other three types to list formal procedure and setting of an example as key motivators. The power-based and discretion-based styles are similar to each other in their stated intentions, and both place a higher emphasis on formal procedure than the affiliative type. In Appendix D, Table A.5 provides a breakdown of the average prevalence of these intentions by vignette.

¹⁴ To generate summary statistics of types, we use the point estimates of $\hat{\theta}_i$ to assign each individual to a model “pure type”, to then assign them to a pure type T_i using: $T_i = \arg \max \theta_i$. Formal inference uses the full distribution of the posterior distribution of θ to account for the issues with two-stage inference highlighted in Battaglia, Christensen, Hansen, and Sacher (2024).

Table 2: Reported intentions by type

	Expresses distrust	Mentions maintaining relationship	Mentions following procedure	Mentions setting example
Rule-based	.300	.259	.433	.378
Affiliative	.200	.332	.290	.298
Power-based	.274	.233	.358	.287
Discretion-based	.257	.218	.321	.285
<i>p</i> -value	<0.001	<0.001	<0.001	<0.001

Notes: This table reports respondents' *ex-ante* intentions, elicited through open-ended questions before they answer each vignette. For each person we record whether they (i) say they do not trust the other party, wish to maintain a good relationship, mention following procedure, and aim to set an example for future interactions. After estimating our model we assign each respondent to the *modal* latent type – the topic k with the highest posterior weight $\hat{\theta}_{ik}$. The reported means are the sample means of the four intention indicators, conditional on latent type. To test whether intentions differ across types, we compute the Mahalanobis (Wald) statistic and compare it to a χ^2_3 distribution. This relies on the multivariate normal approximation to the posterior described in Section 4.1 of Gelman et al. (2013). The resulting *p*-value appears in the final row. Uncertainty from the first-stage LDA estimation is thus propagated via posterior draws.

4 Firms

4.1 Managers’ preferences over management styles

Do these differences reflect alternative approaches to management that may be valued differently by different kinds of firms (as in [Bandiera et al. \(2020\)](#)), or is one type clearly preferred by firms? In this section, we turn to the results from our incentivised elicitation of firm preferences over the vignette responses of our respondents. In Appendix Table [A.25](#), we show that the probability that any two managers agree in their assessment of the same pair of candidates is about 60% (and that this figure is quite stable across vignettes, in both the wage employment and self-employment domains). In Appendix Table [A.26](#), we show that this probability is also largely invariant to similarities between the assessing firms.

Our Augmented Latent Dirichlet Allocation describes K types of studio responses. We now want to estimate firms’ preferences over these styles. We elicited firms’ preferences in the form of rankings over sets of three responses; we therefore estimate using a Plackett-Luce model ([Plackett, 1975](#); [Luce, 1959](#)). This model (sometimes known as the ‘rank-ordered logit’) is specified as follows. For firm f , the latent utility of choosing each candidate i in vignette v is:

$$U_{fiv} = \beta_f \theta_i + \gamma_i + \varepsilon_{fiv}, \quad (2)$$

where θ_i is the vector of type probabilities from the Latent Dirichlet model,¹⁵ and γ_i acts as a candidate-level random effect. We provide more detail on this model in Appendix [E](#).

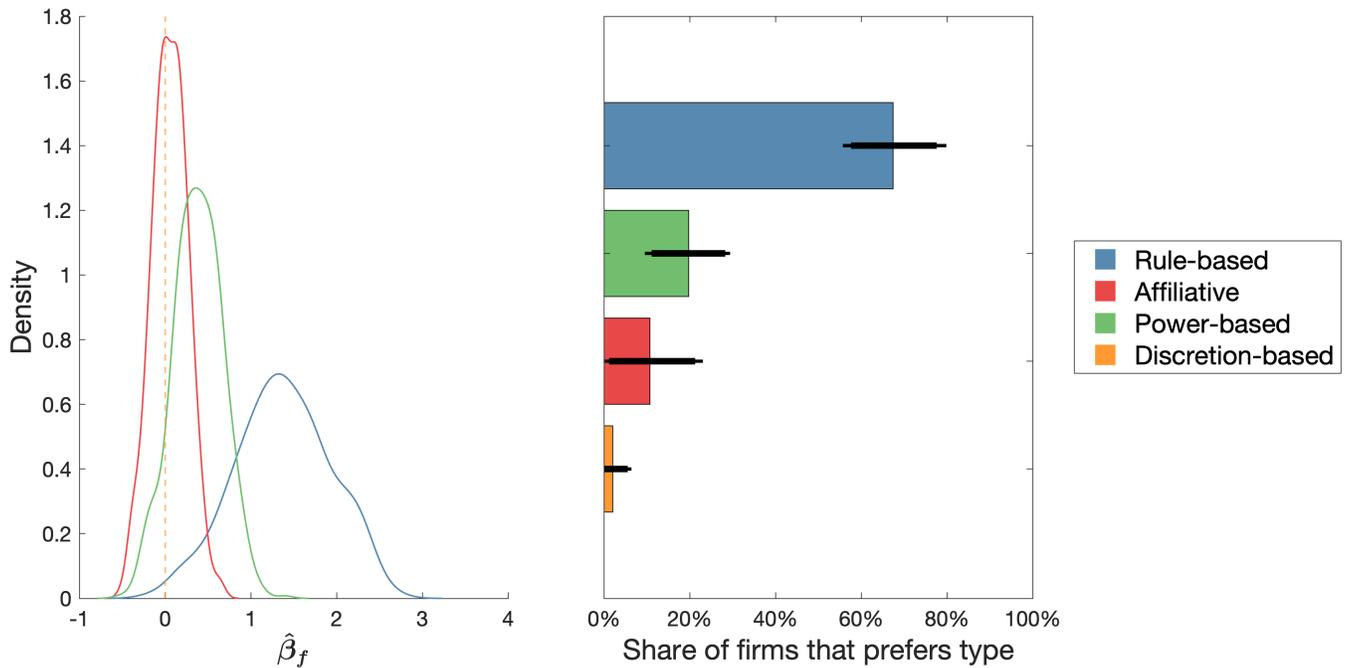
We plot our estimates in Figure [5](#). The hierarchical model generates firm-level estimates for $\hat{\beta}_{fs}$, and plot these estimates in the left panel (where the discretion-based style is used as the omitted category); in the right panel, we show the estimated share of firms that first-preference each type. We find that the rule-based type is strongly preferred by most managers: indeed, about 70% of firms would first-preference that type. About 20% of firms instead first-preference the power-based type. Hiring managers have broadly similar preferences for the discretion-based and affiliative styles, and very few firms would first-preference either type.

To interpret these results, we consider what would happen if we shift the value of θ to represent a ‘pure’ type manager for one member of each triplet. Specifically, for the simulation of the parameter $\beta_{\text{rule-based}}$, we adjust one individual in each triplet to have $\theta_{\text{rule-based},i} = 0.85$, setting the other elements of the vector θ_i to 0.05.¹⁶ We then simulate the utilities of the HR managers using this updated θ and calculate the

¹⁵ We jointly estimate the Latent Dirichlet model and the Plackett-Luce model. Joint estimation avoids the “attenuation” bias that would arise if we treated the type probabilities generated by the Latent Dirichlet model as observed data, which would result in both biased parameters and incorrect standard errors. Following [Battaglia et al. \(2024\)](#), this approach propagates the uncertainty in the Dirichlet parameters θ_i and yields valid inference. See Appendix [E](#) for details, where we also show our estimates for ϕ and θ are not affected by this joint estimation procedure.

¹⁶ This represents approximately the 95th percentile in the estimated distribution of $\hat{\theta}$.

Figure 5: The distribution of the preferences for entry-level managers



Notes The left panel of this figure shows the distribution of the estimates for $\hat{\beta}_f$ across all firms in terms of demand for entry-level managers. In this figure, the discretion-based type is omitted and the estimates for $\hat{\beta}_{\text{Rule-based}}$, $\hat{\beta}_{\text{Affiliative}}$ and $\hat{\beta}_{\text{Power-based}}$ are plotted. The right panel shows that share of firms that most values a specific managerial type across our draws, i.e. the share of firms for which, for example, $\hat{\beta}_{\text{Rule-based}}$ is the largest element of $\hat{\beta}_f$ for that draw. We calculate both the mean, and 90 and 95% credible sets based on our draws.

updated probability that this simulated individual is the winner. Initially, the probability each individual wins pairwise comparisons is 50%. We find this probability increases to 62% for the rule-based type, decreases to 45% for the affiliative and discretion-based types, and slightly decreases to just under 50% for the power-based type. Another way of thinking of this is through a variance decomposition – which shows that there is substantial idiosyncratic noise, but that over half of the non-idiosyncratic variation in preferences can be explained by the preferences over our four types.¹⁷

In Appendix Figure A.11, we examine the distribution of the parameter estimates for β_f for subsets of the firms. We split the sample (i) by sector (manufacturing versus services), (ii) by whether the manager has an above- or below-median score on the World Management Survey, and (iii) by whether the firm is public or private. We find no meaningful relationship between the distribution of β_f and any of these different dimensions.

We then estimate two extensions of the model, both offering greater flexibility than the main specification. First, in Appendix Table A.21, we allow managers’ preferences over types to vary at the vignette level (that is, we estimate β_{fv} rather than β_f). We find broadly similar preferences across all vignettes – in particular, a strong preference for the rule-based type in each of the vignettes – with some interesting heterogeneity in preferences for the power-based type (firms are more receptive to the power-based style in the scenarios where the firm clearly holds a dominant role in the relationship – namely, in vignettes featuring employee absence, a request for a pay rise, and the request from a client). Second, in Appendix Table A.22, we allow managers’ preferences over types to vary based on the gender of the young professional. We find no significant difference in preferences by candidate gender (the credible sets on the interaction include zero in all cases); looking solely at point estimates, we note that, if anything, the rule-based management style is particularly valued among female candidates.

¹⁷ Specifically, we decompose the variance of the parameter estimates of the model specified in Equation 2. We calculate: (a) the variance of the total non-random component of the utility function, (b) the variance of just the $\beta_f \theta_i$ term and (c) the variance of γ_i . We obtain $\widehat{\text{var}}(\beta_f \theta_i + \gamma_i) = 0.63$, $\widehat{\text{var}}(\beta_f \theta_i) = 0.34$ and $\widehat{\text{var}}(\gamma_i) = 0.29$.

4.2 Robustness: Random variation in actors' gender

We previously described the five vignettes that were played to young professionals. As part of this experimental design, we additionally randomized both (i) the order in which young professionals viewed the five vignettes and (ii) the gender of the actor shown in each vignette (randomised independently across vignettes, having recorded each vignette twice: once with a male actor and once with a female actor). In this final part of the analysis, we exploit this additional exogenous variation to generate a robustness check on our main results. Specifically, we exploit the random variation in whether the first actor viewed by a respondent is female.¹⁸

To analyse this variation, we run three kinds of analysis – all reported in Table 3. First, we test the effect upon managerial types of having seen a female actor in the first vignette. To do this, we use the Hamiltonian Monte Carlo chains, and compare the posterior distributions between those respondents who initially saw a male actor and those who initially saw a female actor; we then use these distributions to construct point estimates of the causal effect, and 95% Bayesian credible intervals. Panel A shows that exposure to a female actor in the first vignette leads respondents to rely more on the rule-based management style and less on the power-based style (with a 95% credible interval that excludes zero for both effects).

In Panel B, we report a series of OLS regressions to test the average treatment effect on measured attributes of professionals' responses; we estimate:

$$Y_{ive} = \alpha + \beta \cdot \text{First_Actor_Female}_i + \gamma_v + \lambda_e + \varepsilon_{ive}, \quad (3)$$

where we include vignette (γ_v) and enumerator (λ_e) fixed effects, and cluster errors at the individual (i) level. We examine a set of four outcome variables: a binary indicator for relying on the firm's interest as justification; an indicator relying on the respondent's own interest as justification, an indicator for relying on a formal policy as source of authority, and an indicator for relying on seniority as source of authority.¹⁹ We restrict the sample to data from the second through fifth vignettes a respondent sees – to exclude any direct effects through linguistic gender clues. We find that, where the first actor is a woman, respondents are more likely to rely on the firm's interest instead of their own interest as their justification, and more likely to rely on formal authority instead of seniority. This confirms the result in Panel A: that, on average,

¹⁸ In principle, we could also conduct a similar analysis using the gender of each separate actor – with the analysis proceeding at the level of the respondent-vignette pair. Given our particular focus on managerial types – that is, a property that we estimate at the level of the respondent – it makes more sense to use this simpler respondent-level variation.

¹⁹ These specific variables are selected to illustrate the shift from a power-based to a rule-based management style, we report a series of multinomial logit regressions of the effect of seeing a female actor on the full set of attributes from which these were selected in Appendix L.

respondents exhibit more “rule-based” behaviour after being exposed to a female actor in the first vignette.

Table 3: Actors’ gender and management styles

Panel A: Effect on estimated types				
	Rule-based	Affiliative	Power-based	Discretion-based
First actor female	0.030***	-0.004	-0.030***	0.003
Constant	0.262	0.191	0.296	0.251
Bayesian Credible Interval	[.012 .049]	[-.021 .013]	[-.051 -.006]	[-.021 .031]
N	982	982	982	982
Panel B: Effect on selected attributes				
	Justification		Authority	
	Rely on firm’s interest	Rely on own interest	Rely on formal authority	Rely on seniority
First actor female	0.032** (0.015)	-0.023** (0.011)	0.021* (0.011)	-0.026 (0.016)
Constant	0.547*** (0.011)	0.177*** (0.008)	0.121*** (0.008)	0.507*** (0.012)
Enumerator FE	Yes	Yes	Yes	Yes
Vignette FE	Yes	Yes	Yes	Yes
Mean dep. var	0.547	0.170	0.132	0.488
N	7860	7860	7860	7860
Panel C: Effect on managers’ assessments				
	Ranking Data		Normalised likert score	
	Manager	Entrepreneur	Manager	Entrepreneur
First actor female	0.052** [0.010 0.093]	0.060*** [0.017 0.102]	0.083** (0.033)	0.093*** (0.033)
Constant	0.474*** [0.453 0.493]	0.470*** [0.449 0.492]	-0.045* (0.025)	-0.050** (0.025)
Vignette FE			Yes	Yes
N			6874	6869

Notes: This figure displays the causal relationship between the first actor a respondent sees and their subsequent responses (for the second to fifth vignette). Panel A shows a causal effect on adopting a more rule-based management style (with 95% Bayesian credible intervals in square brackets). In Panel B we report the effect of seeing a female actor on selected attributes exhibited by the respondent in subsequent vignettes based on a linear regression. These attributes are selected to illustrate the shift from a power-based to a rule-based management style; Appendix L reports a sequence of multinomial logit regressions on the full set of attributes from which these were selected. The third panel shows the effect of the first actor being female on the HR managers’ rankings. The constant is the probability of winning each pairwise comparison after seeing a male actor in the first vignette, with the 95% credible set reported. The parameter “First actor female” is the coefficient β in probability space, i.e. the mean marginal effect, again with the 95% credible set in square brackets. The latter two columns use the non-incentivised, normalised Likert score that the HR managers give each candidate. We implement a simple OLS regression with vignette fixed effects with standard errors clustered at the HR manager level. Statistical significance (in the classical sense for regressions, and testing whether the 90, 95 or 99% credible set contain zero for Bayesian analysis), is denoted where appropriate by * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

Finally, we test the impact upon firms’ responses – stressing that the HR managers were never told whether the young professionals had seen a male or female actor for their first vignette.²⁰ To assess this, we estimate a Plackett-Luce model with a single covariate; we define the latent utility of firm f ’s assessment of candidate i in vignette v as:

$$y_{fiv}^* = \beta \cdot \text{First_Actor_Female}_i + \varepsilon_{fiv}. \quad (4)$$

Panel C of Table 3 reports the results from this exercise. Specifically, we calculate the implied probability that an HR manager ranks a candidate who first saw a female actor more highly than one who first saw a male actor. We estimate that seeing a female actor first increases by about 5.2 percentage points the probability that an HR manager would prefer the candidate as a prospective entry-level manager, and by about 6 percentage points the probability that an HR manager would prefer the candidate as a prospective entrepreneur; the credible sets for each of these estimates excludes zero.

In sum, the randomisation of actors’ gender generates additional excludable variation in respondents’ measured managerial traits – and we show that, consistently with our main results, firms prefer managers exhibiting the ‘rule-based’ style.

4.3 Workers’ preferences for management styles: A model-informed field survey

The previous subsections show a clear preference from firm managers for the rule-based management style. To what extent is this preference shared by firm employees? On the one hand, for example, employees might share managers’ preference for the rule-based management style– perhaps because they respect the way that this type communicates clearly on the basis of set rules. On the other hand, employees might prefer the more amiable interpersonal style of the affiliative type.

To answer this question, we returned to the field for a series of semi-structured interviews with firm employees. We did this after obtaining our model estimates – so that the interviews could focus on interpreting our model results. Specifically, we interviewed a total of 66 employees, drawn from 17 firms in our original sample; we did this through a series of 12 focus groups in Addis Ababa, in May 2025. The average worker was 31 years old and had spent four years with their firm, which were mostly operating in manufacturing (51.5 percent) and services (42.4 percent). Around 38% of these respondents had completed at least a bachelor’s degree, compared to 78% of our young professionals. Each respondent viewed two vignettes: the vignette concerning line-management of an employee, and the vignette

²⁰ Amharic marks second-person agreement in both pronouns and verb conjugation – using masculine or feminine forms for gendered addressees, plus a gender-neutral honorific. In the videos of our actors, speakers almost always used the gender-neutral honorific form (in nine of our ten videos, the speaker uses the honorific; in the tenth video, the speaker begins with a gendered pronoun and matching verb form before switching to the honorific). However, many of the respondents used the gendered pronoun and matching verb form; for this reason, the exclusion restriction is more difficult to justify for the first vignette.

concerning a requested pay rise. (We chose these two vignettes for their direct relationship to employee management.) For each of these two vignettes the respondent then watched the responses of a different set of five young professionals (each of whom had given permission for their video to be viewed in this way). These five young professionals were chosen based on the model estimates; specifically, we chose (i) a professional close to being a pure rule-based type, (ii) a professional close to being a pure affiliative type, (iii) a professional close to being a pure power-based type, (iv) a professional close to being a pure discretion-based type, and (v) a professional chosen for being approximately an equal weighting of those four types. (That is, in the context of the tetrahedron illustrated in Figure 3, we chose young professionals from near to each of the four vertices and from close to the centroid; we illustrate this in Appendix Figure A.12.) We used a total of 60 young professionals (that is, 12 in each of these five categories), to ensure that our interview results are not driven by the idiosyncracies of specific individual videos. Further, we ensured that no employee saw the same professional more than once, and that each employee either saw only male professionals or only female professionals. We used neutral labels to refer to all of the videos.

We then asked each employee a series of questions about the suitability of each young professional for their firm. Specifically, we asked each employee to provide a Likert-scale assessment of (i) whether the young professional would enhance firm productivity, (ii) whether the young professional would help the firm to achieve its targets, (iii) whether the employee would be happy to work hard for the young professional, (iv) whether the young professional would treat different employees equally, and (v) whether the style of response is typical for the employee's firm. Finally, for both vignettes, we asked each employee to rank the five professionals they had seen in order of their likely contribution to the long-term success of the firm.

Figure 6 shows the distribution of these overall rankings. The figure shows a clear and significant preference for professionals of the rule-based managerial style. In Appendix Table A.29, we find the same clear pattern across all of the questions asked: the professionals with a rule-based style were significantly preferred to all other professionals across each of the questions asked. Qualitative follow-up responses placed particular weight on employees' perceptions that they would be treated fairly and consistently; Appendix Table A.29 shows a clear preference for the rule-based style in this regard, and the qualitative discussion suggests that this was a central reason for the strong overall assessment of professionals with this managerial style. Indeed, in this regard, it is interesting to note that the affiliative professionals were the least preferred type across all questions asked – including in employees' perceptions that such professionals would treat employees equally. Qualitative follow-up discussion suggested that – far from seeing this type as being their ally or supporter – employees may have viewed affiliative professionals as being overly impressionable and prone to favouritism. In Appendix Tables A.30 and A.31, we disaggregate by vignette; we find that these patterns are very similar across the two vignettes (with the professionals with the discretion-based managerial style performing relatively better in the line management scenario than in the pay rise scenario).

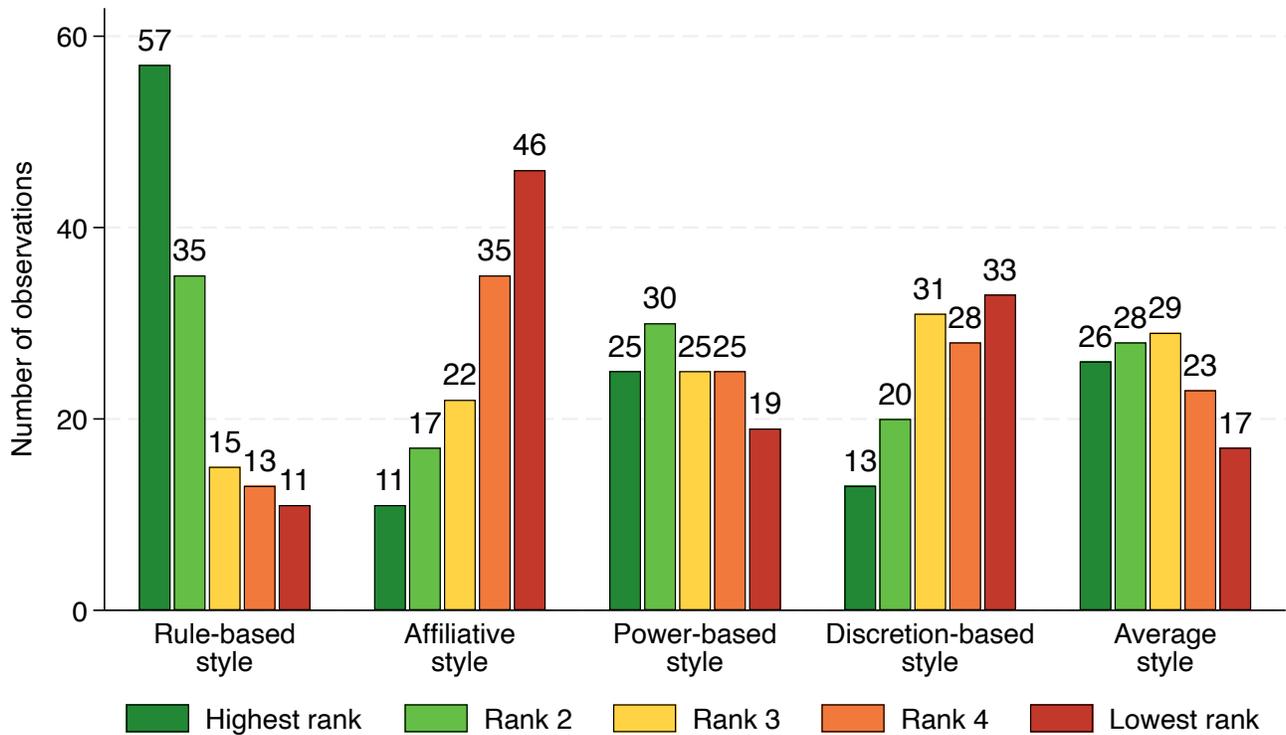


Figure 6: Distribution of ranks by management style

Notes: This figure summarizes 132 rankings of 66 workers of the effectiveness of our young professionals as managers. For each “pure” type, we count how many times that type is ranked highest, second, third, fourth, or fifth and display these counts here. We conduct all statistical inference for this section in Appendix Table A.29 using a rank-ordered logit model. We note that the rule-based managerial style is significantly preferred over all other types ($p \leq 0.001$ in a binary comparison of the coefficients). Next, the power-based and average type are similarly preferred over the affiliative, and discretion-based style, and finally the discretion-based type is preferred over the affiliative type.

5 Experimental variation in managerial exposure

5.1 The causal effects of managerial experience

The results in section 4 show that rule-based managers are preferred by firms; the earlier results in section 3 show descriptively that rule-based managers also have historically greater labour market attachment and higher parental education. Together, these results suggest that managerial traits may be a mechanism for persistence in labour market exclusion – but do not tell us whether labour market experience influences the development of managerial traits, or whether managerial traits, instead, determine labour market outcomes. To answer this question, we turn to evidence from respondents’ participation in a management experience experiment five years prior to attending the studio.

For each style k , let $\theta_{ik}^{(d)}$ denote the latent trait for individual i at Hamiltonian Monte Carlo (HMC) draw d . We define the draw-specific group mean within treatment arm $t \in \{\text{treated}, \text{control}\}$ as

$$\bar{\theta}_{k,t}^{(d)} = \frac{1}{n_t} \sum_{i: T_i=t} \theta_{ik}^{(d)},$$

where T_i indicates the treatment assignment of individual i and n_t is the number of individuals in arm t .

We compute the (draw-specific) average treatment effect as

$$\Delta\theta_k^{(d)} = \bar{\theta}_{k,\text{treated}}^{(d)} - \bar{\theta}_{k,\text{control}}^{(d)}.$$

For each style k , we report the posterior mean and 95% credible interval of the set $\{\Delta\theta_k^{(d)}\}_d$. We conduct this exercise for three groups: the full sample, individuals whose parents did not complete primary school, and those with at least one parent who did.

We find, in Table 4, that the management experience treatment makes individuals on average more rule-based and less discretion-based (column 2), and that this result is entirely driven by individuals whose parents have not finished primary school (column 4). We find no evidence for any treatment effect for individuals whose parents finished primary school (column 6).

We draw two conclusions from this result. First, management experience shapes an individual’s management style. Treated individuals became significantly more likely to display a rule-based style of management, and less likely to display a discretion-based style. Since the treatment focused on shadowing a manager, this suggests that the intervention shows what effective management looks like rather working through “learning by doing”. Second, we find that this average treatment effect is *entirely* driven by participants for whom neither parent completed primary school. We interpret this finding as showing that these individuals do not learn the established norms from their parents, as shown by the lower weight on the rule-based management style in the absence of the treatment, but are able to learn these norms through

exposure to formal management practices at an established firm. Strikingly, the treatment significantly reduces the average difference in management styles between individuals for whom at least one parent completed primary school and for those for whom neither did so - from 8.0 percentage points for untreated individuals to 2.3 percentage points for treated individuals. These findings suggest that managerial experience plays an important role in shaping managerial traits, particularly for individuals with less prior exposure to formal firms.

Table 4: The causal effect of managerial experience on management style by parents education

	Full sample		Low parental education		High parental education	
	(1)	(2)	(3)	(4)	(5)	(6)
Rule-based (%)	26.5	2.4	22.2	5.6	30.2	-0.1
		[0.2, 4.5]		[2.4, 8.9]		[-3.2, 2.8]
Affiliative (%)	18.4	1.2	19.3	0.7	17.7	1.5
		[-0.5, 2.9]		[-2.0, 3.3]		[-1.1, 4.1]
Power-based (%)	28.0	-0.1	28.7	-0.8	27.1	0.9
		[-2.3, 2.3]		[-4.5, 2.7]		[-2.5, 4.2]
Discretion-based (%)	27.1	-3.6	29.7	-5.6	24.9	-2.3
		[-5.8, -1.5]		[-8.6, -2.1]		[-5.5, 0.9]
N	479	500	222	237	250	259

Notes This table reports the treatment effect of the management experience experiment on the managerial traits of individuals. The treatment effect is calculated based on the distribution of the difference in the average value of θ for treated and untreated individuals. Columns (1), (3) and (5) report the average estimated value of θ_i for individuals that were not treated in the management experience experiment for respectively all individuals, individuals whose parents did not finish primary school and for individuals for whom at least one parent did. Columns (2), (4) and (6) report the treatment effect of the management experience experiment on their managerial traits for these three groups respectively. In columns (2), (4) and (6) both the average treatment effect and the 95% credible interval, in square brackets, are reported. For individuals where we only record the education of one parent we use this to classify individuals as low and high parental education. We drop individuals for whom we observe neither parents' education level from the heterogeneity analysis.

5.2 The causal effect on managers' assessments

We next test whether this experimental managerial exposure affects how candidates are ranked by HR managers. We estimate a Plackett–Luce (rank-ordered logit) model on firms' ranking data and translate posterior draws into pairwise winning probabilities. For firm f , the latent utility of choosing each candidate i in vignette v is:

$$U_{fiv} = \delta_1 \cdot T_i + \delta_2 \cdot \text{PE}_i + \delta_3 \cdot T_i \cdot \text{PE}_i + \varepsilon_{fiv}, \quad (5)$$

where T_i is a dummy equal to one if the individual is treated, PE is a dummy equal to one if either parent has completed primary school and ε_{fiv} is a Type-1 EV distributed error term. We estimate this model in Stan.

To interpret these results, we use the posterior draws to estimate winning probabilities for various pairwise group comparisons, keeping fixed either parental education or treatment status. In Table 5 we first show that, across the full sample, a treated individual has a 53.3% chance of winning in a binary comparison with an untreated individual – suggesting the effect of the treatment on managerial styles is reflected in HR managers' assessments. Next, we examine the treatment effect conditional on parental education. For individuals whose parents did not complete primary school, a treated individual has a 55.1% chance of winning when compared to an untreated individual; for individuals whose parents did complete primary school this is equal to 51.3%, which is not estimated to be significantly different than 50%. Finally, we show that, in the absence of treatment, individuals with low parental education are at a disadvantage; however, conditional on treatment, parental education no longer affects managers' assessments.

Table 5: The causal effect of managerial experience on managers' assessments

Comparison	Winning Probability (%)
<i>Overall Treatment Effect</i>	
Treated vs. Untreated (All)	53.3*** [51.6, 55.0]
<i>By Parental Education</i>	
Treated vs. Untreated (Low parental education)	55.1*** [52.6, 57.4]
Treated vs. Untreated (High parental education)	51.3 [49.0, 53.7]
<i>Gap in Absence of Treatment</i>	
Low vs. High Parental Education (Untreated)	45.3*** [42.8, 47.7]
<i>Gap Conditional On Being Treated</i>	
Low vs. High Parental Education (Treated)	49.7 [46.3, 51.0]

Notes: Entries report posterior means of pairwise “winning” probabilities from a Plackett–Luce (rank-ordered logit) model. For any two profiles i and j , the winning probability is $\Pr(i \succ j) = \frac{\exp(U_{fiv})}{\exp(U_{fiv}) + \exp(U_{fjv})}$, computed at each posterior draw and then averaged. Brackets show the 95% credible intervals calculated as the range between the 2.5th and 97.5th percentile of the distribution of the posterior draws. We show the heterogeneity of this treatment effect by parental education, and the remaining gap by parental education conditional on treatment status. *Treated vs. Untreated (All)* reports the treatment effect for the full sample. *By Parental Education* conditions on low vs. high parental education (low means neither parent completed primary school; high mean at least one parent completed primary school). *Low vs. High Parental Education (Untreated/Treated)* compares otherwise identical profiles that differ only in parental education, holding treatment status fixed. A value of 50% indicates no difference in expected rank in a binary comparison. Appendix Table A.24 reports the underlying coefficient estimates in log-odds space.

6 Discussion

In this paper, we have developed a novel experimental approach to assessing managerial traits – using a controlled studio setting to observe and analyse how young Ethiopian professionals respond to realistic workplace challenges. By combining experimental data with Bayesian hierarchical modeling, we identified four distinct managerial styles, which we label as ‘rule-based’, ‘affiliative’, ‘power-based’ and ‘discretion-based’. This exercise has generated several key insights.

First, we show meaningful heterogeneity in managerial traits: our four archetypes reflect meaningful variation in how young professionals approach decision-making and conflict resolution in managerial roles. The rule-based style, for instance, emerges as the most dominant managerial style, particularly valued by firms for entry-level managerial roles. This type tends to emphasise formal procedures and the firms’ interests – contrasting sharply, for example, with the affiliative type – who seeks shared ground and often concedes to a counterparty.

Second, our analysis reveals that firms in Ethiopia, across a range of industries, tend to prefer a rule-based management style for entry-level positions. Some previous studies (e.g. [Bandiera et al., 2020](#)) have emphasized the heterogeneity in firm preferences for senior managers; our results suggest that firms may place greater value on uniformity in managerial traits at the lower levels of management. This homogeneity could reflect the need for clear, consistent, and rule-based decision-making among less-experienced managers. This result accords with several anecdotal characterisations of large Ethiopian firms – which tend to emphasise the importance of strong centralised leadership, with a clear role for corporate hierarchy.

Third, we show that young professionals with more exposure to the labor market are substantially more likely to exhibit traits of the rule-based archetype. Our evidence based on random assignment to a previous management placement experiment suggests this is at least partially a causal link. This finding suggests that managerial styles are not purely innate, but can be shaped and developed through professional experience. These results align with previous studies showing the importance of early exposure to managerial environments in shaping long-term behavioural traits ([Bertrand and Schoar, 2003](#); [Malmendier et al., 2011](#)).

Methodologically, our paper demonstrates the feasibility and the value of using studio-based vignettes to assess managerial traits. Unlike traditional observational studies, our controlled setting allows us to systematically compare managerial responses across a fixed set of scenarios, removing the endogeneity inherent in real-world managerial decision-making. This approach opens up new avenues for studying managerial behavior – particularly in contexts where direct observation of managerial actions may be impractical or biased by the organizational setting.

More generally, our study complements the growing literature on labor market exclusion faced by young professionals. The finding that managers with a rule-based style are both in higher demand by firms and more likely to have greater labor market exposure suggests that managerial traits may serve as a key

mechanism by which some young professionals achieve sustained labor market success, while others face exclusion. Our results on the relevance of parental education speak particularly to this issue. Respondents with less educated parents are significantly less likely to exhibit a rule-based style than respondents with more educated parents, and are significantly less likely to be preferred by firms (notwithstanding that firms have no direct means of observing respondents' parental education). However, the exogenous variation in labour market exposure essentially closes completely both of these gaps.

This suggests several novel angles for policy innovation. Policy to support disadvantaged jobseekers has traditionally focused on development of human capital, or on effective signalling of existing skills (Abebe, Caria, Fafchamps, Falco, Franklin, and Quinn, 2021b). Our paper suggests a third possibility: working with young professionals to develop the kind of rule-based professional traits sought by large firms. Our results show that spending time in established firms is one mechanism to achieve this; there may be several others.²¹

Future research could expand on this work by exploring how these preferences evolve as managers move into higher-level roles and by investigating the role of cultural and institutional factors in shaping managerial behaviour. Additionally, the scalability of our experimental method presents opportunities for further research across different contexts and labour markets, providing a useful tool for understanding the development of managerial talent in a wide range of settings.

²¹ In a related space, several initiatives seek to support disadvantaged jobseekers by targeting interview appearance, among other skills. For example, 'Dress for Success Worldwide' provides 'professional attire services' as part of its support for women jobseekers; in Ethiopia, 'The Talent Firm' is a professional recruitment company that provides 'The Professional's Closet': a service providing donated professional attire for professional jobseekers.

References

- Alberto Alesina, Sebastian Hohmann, Stelios Michalopoulos, and Elias Papaioannou. Intergenerational mobility in africa. *Econometrica*, 89(1):1–35, 2021. 10.3982/ECTA17018.
- Sam Asher, Paul Novosad, and Charlie Rafkin. Intergenerational mobility in education in india. *American Economic Journal: Applied Economics*, 14(2):213–241, 2022. 10.1257/app.20200389.
- Girum Abebe, A. Stefano Caria, and Esteban Ortiz-Ospina. The selection of talent: Experimental and structural evidence from ethiopia. *American Economic Review*, 111(6):1757–1806, June 2021a. 10.1257/aer.20190586.
- Eliana Carranza, Robert Garlick, Kate Orkin, and Neil Rankin. Job search and hiring with limited information about workseekers’ skills. *American Economic Review*, 112(11):3547–3583, November 2022. 10.1257/aer.20200961.
- Samuel Bowles and Herbert Gintis. Schooling and the abandonment of the socialization hypothesis in economic theory. *The Review of Economics and Statistics*, 59(4):487–501, 1977. 10.2307/1928709.
- Samuel Bowles and Herbert Gintis. Schooling in capitalist america revisited. *Scandinavian Journal of Economics*, 104(1):1–24, 2002. 10.1111/1467-9442.00267.
- James J. Heckman and Yona Rubinstein. The importance of noncognitive skills: Lessons from the ged testing program. *American Economic Review*, 91(2):145–149, 2001. 10.1257/aer.91.2.145.
- Oriana Bandiera, Stephen Hansen, Andrea Prat, and Raffaella Sadun. Ceo behavior and firm performance. *Journal of Political Economy*, 128(4):1325–1369, 2020.
- Amanda Dahlstrand, Dávid László, Helena Schweiger, Oriana Bandiera, Andrea Prat, and Raffaella Sadun. Ceo-firm matches and productivity in 42 countries. Technical report, National Bureau of Economic Research, 2025.
- Girum Abebe, Marcel Fafchamps, Michael Koelle, and Simon Quinn. Matching, management and employment outcomes: A field experiment with firm internships. *Working paper*, 2024.
- D Goleman. Leadership that gets results. *Harvard Business Review*, 78(2):78–89, 2000.
- Alan M Benson and Kathryn L Shaw. What do managers do? an economist’s perspective. *NBER Working Paper 33431*, 2025.
- David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022, 2003.
- Thomas L Griffiths. Finding scientific topics. *PNAS*, 2004.
- Vittorio Bassi and Aisha Nansamba. Screening and signalling non-cognitive skills: experimental evidence from uganda. *The Economic Journal*, 132(642):471–511, 2022.
- Pierre Bourdieu. The forms of capital. In John G. Richardson, editor, *Handbook of Theory and Research for the Sociology of Education*, pages 241–258. Greenwood, New York, 1986.
- Seth D. Zimmerman. Elite colleges and upward mobility to top jobs and top incomes. *American Eco-*

-
- nomics Review*, 109(1):1–47, 2019. 10.1257/aer.20171019. URL <https://www.aeaweb.org/articles?id=10.1257/aer.20171019>.
- Armin Falk, Fabian Kosse, and Pia Pinger. Mentoring and schooling decisions: Causal evidence. (13387), 2020. URL <https://www.iza.org/publications/dp/13387/mentoring-and-schooling-decisions-causal-evidence>.
- Andrés Barrios-Fernández, Christopher Neilson, and Seth Zimmerman. Elite universities and the intergenerational transmission of human and social capital. Working paper, August 2024. URL https://andresbarriosf.github.io/igt_hcsc.pdf.
- Uditi Karna, John List, Andrew Simon, and Haruka Uchida. Five facts about the first-generation excellence gap. Technical report, The Field Experiments Website, 2025.
- Soumitra Shukla. Making the elite: Class discrimination at top firms. Working paper, 2025. URL https://soumitrashukla.github.io/making_the_elite_shukla.pdf.
- Orazio P Attanasio, Áureo De Paula, and Alessandro Toppeta. Intergenerational mobility in socio-emotional skills. *Journal of Public Economics*, 248:105423, 2025.
- Daniela Scur, Raffaella Sadun, John Van Reenen, Renata Lemos, and Nicholas Bloom. The world management survey at 18: lessons and the way forward. *Oxford Review of Economic Policy*, 37(2): 231–258, 2021.
- Marianne Bertrand and Antoinette Schoar. Managing with style: The effect of managers on firm policies. *The Quarterly Journal of Economics*, 118(4):1169–1208, 2003.
- Nicholas Bloom, Erik Brynjolfsson, Lucia Foster, Ron S. Jarmin, Itay Patnaik, Itzhak Saporta-Eksten, and John Van Reenen. What drives differences in management practices? *American Economic Review*, 109(5):1648–1683, 2019. 10.1257/aer.20170491.
- Mitchell Hoffman and Steven Tadelis. People management skills, employee attrition, and manager rewards: An empirical analysis. *Journal of political economy*, 129(1):243–285, 2021.
- Edward P Lazear, Kathryn L Shaw, and Christopher T Stanton. The value of bosses. *Journal of Labor Economics*, 33(4):823–861, 2015.
- Robert D Metcalfe, Alexandre B Sollaci, and Chad Syverson. Managers and productivity in retail. Technical report, National Bureau of Economic Research, 2023.
- Ritwika Sen. Supervision at work: Evidence from a field experiment. Technical report, Working paper, 2024.
- Jonas Hjort, Hannes Malmborg, and Todd Schoellman. The missing middle managers: Labor costs, firm structure, and development. Technical report, National Bureau of Economic Research, 2022.
- Ulrike Malmendier, Geoffrey Tate, and Jon Yan. Overconfidence and early-life experiences: The effect of managerial traits on corporate financial policies. *The Journal of Finance*, 66(5):1687–1733, 2011.
- Steven N Kaplan, Mark M Klebanov, and Morten Sorensen. Which CEO characteristics and abilities matter? *The Journal of Finance*, 67(3):973–1007, 2012.

-
- Oriana Bandiera, Luigi Guiso, Andrea Prat, and Raffaella Sadun. Matching firms, managers, and incentives. *Journal of Labor Economics*, 33(3):623–681, 2015.
- Efraim Benmelech and Carola Frydman. Military ceos. *Journal of Financial Economics*, 117(1):43–59, 2015.
- Paula López-Peña, Muntasir Mozumder, Abdul Rabbani, and Christopher Woodruff. Toxic Managers, Firm Productivity, and Worker Well-Being: Evidence from Bangladeshi Garment Factories. *CEPR Discussion Paper No. 19936.*, 2025. URL <https://cepr.org/publications/dp19936>.
- Ben Weidmann, Joseph Vecci, Farah Said, David J. Deming, and Sonia R. Bhalotra. How do you find a good manager? Working Paper 32699, National Bureau of Economic Research, July 2024. URL <https://www.nber.org/papers/w32699>.
- Xiangyu Chang, Lili Dai, Lingbing Feng, Jianlei Han, Jing Shi, and Bohui Zhang. A good sketch is better than a long speech: evaluate delinquency risk through real-time video analysis. *Review of Finance*, 29(2):467–500, 2025.
- Amy Handlan and Haoyu Sheng. Gender and tone in recorded economics presentations: Audio analysis with machine learning. *Working paper*, 2023.
- Yuriy Gorodnichenko, Tho Pham, and Oleksandr Talavera. The voice of monetary policy. *American Economic Review*, 113(2):548–584, 2023.
- Marius Guenzel, Shimon Kogan, Marina Niessner, and Kelly Shue. Ai personality extraction from faces: Labor market implications. *SSRN Electronic Journal*, January 2025. 10.2139/ssrn.5089827. URL <https://ssrn.com/abstract=5089827>.
- Jens Ludwig and Sendhil Mullainathan. Machine learning as a tool for hypothesis generation. *The Quarterly Journal of Economics*, 139(2):751–827, 2023. 10.1093/qje/qjad001.
- Mark Borgschulte, Marius Guenzel, Canyao Liu, and Ulrike Malmendier. Ceo stress, aging, and death. Working Paper 28550, National Bureau of Economic Research, March 2021. URL <http://www.nber.org/papers/w28550>.
- Stefanie Stantcheva. Understanding tax policy: How do people reason? *The Quarterly Journal of Economics*, 136(4):2309–2369, 2021.
- Ingar Haaland, Christopher Roth, Stefanie Stantcheva, and Johannes Wohlfart. Measuring what is top of mind. *Working paper*, 2024.
- Rocco Macchiavello, Andreas Menzel, Atonu Rabbani, and Christopher Woodruff. Challenges of change: An experiment promoting women to managerial roles in the bangladeshi garment sector. Working Paper 27606, National Bureau of Economic Research, July 2020. URL <http://www.nber.org/papers/w27606>.
- Ben Weidmann and David J. Deming. Team players: How social skills improve team performance. *Econometrica*, 89(6):2637–2657, 2021. 10.3982/ECTA18461. URL <https://doi.org/10.3982/ECTA18461>.

-
- Ben Weidmann, Yixian Xu, and David J Deming. Measuring human leadership skills with ai agents. Technical report, National Bureau of Economic Research, 2025.
- Shengwu Li. Obviously strategy-proof mechanisms. *The American Economic Review*, 107(11):3257–3287, 2017. ISSN 00028282, 19447981. URL <http://www.jstor.org/stable/44871788>.
- Ethiopian Statistical Service and World Bank. Ethiopia socioeconomic panel survey (esps) report - wave 5, 2021/22, September 2023. Published in September 2023.
- Rasmane Ouedraogo and Nicolas Syrichas. *Intergenerational social mobility in Africa since 1920*. International Monetary Fund, 2021.
- Dan Biderman, David Blei, Wei Cai, Andrea Ciccarone, Amir Feder, and Andrea Prat. Fixed effects topic model: Unsupervised identification of organizational culture from job text reviews. in progress. Working paper.
- Stan Development Team. *Stan Reference Manual*, 2024. URL <https://mc-stan.org>. v2.36.0 <https://mc-stan.org>.
- A. Gelman, J.B. Carlin, H.S. Stern, D.B. Dunson, A. Vehtari, and D.B. Rubin. *Bayesian Data Analysis, Third Edition*. Chapman & Hall/CRC Texts in Statistical Science. Taylor & Francis, 2013. ISBN 9781439840955. URL <https://books.google.it/books?id=ZXL6AQAAQBAJ>.
- Stephen Hansen, Raffaella Sadun, Tejas Ramdas, and Joe Fuller. The demand for executive skills. (21-133), 2021.
- Duncan Webb. Silence to solidarity: Using group dynamics to reduce anti-transgender discrimination in india. 2024.
- Laura Battaglia, Timothy Christensen, Stephen Hansen, and Szymon Sacher. Inference for regression with variables generated by ai or machine learning. 2024.
- Robin L. Plackett. The analysis of permutations. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 24(2):193–202, 1975. 10.2307/2346567.
- R. Duncan Luce. *Individual Choice Behavior: A Theoretical Analysis*. John Wiley & Sons, New York, 1959.
- Girum Abebe, Stefano Caria, Marcel Fafchamps, Paolo Falco, Simon Franklin, and Simon Quinn. Anonymity or distance? job search and labour market exclusion in a growing african city. *Review of Economic Studies*, 88(3):1279–1310, 2021b. 10.1093/restud/rdaa057.