UC Davis
UC Davis Electronic Theses and Dissertations

Title

The Economics of University Peer Effects and Employee Training in Microenterprises in Uganda

Permalink

https://escholarship.org/uc/item/3vm4k7p4

Author

Ahimbisibwe, Isaac

Publication Date

2024

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

The Economics of University Peer Effects and Employee Training in Microenterprises in Uganda

By

ISAAC AHIMBISIBWE
DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Agricultural and Resource Economics

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Travis Lybbert, Chair

Scott Carrell

Stephen Boucher

Committee in Charge

2024

Abstract

Human capital development is crucial for economic growth. Policymakers must understand the factors influencing human capital accumulation since they may vary by setting and sector. For example, micro-enterprise employees enhance their skills through on-the-job training or vocational programs. In educational institutions, family and school investments and peer interactions can affect human capital growth. Studies show that characteristics of peers, such as the gender of a classmate, can impact academic performance in elementary school. Other research has highlighted the importance of homophily in social network formation. Other research has highlighted the importance of homophily in social network formation. The latter is especially important in higher education settings, where individuals are likely to engage in assortative matching based on characteristics such as shared identity. Thus, the ethnicity of peers may also play a crucial role in human capital development in settings of ethnic diversity.

The first essay in this dissertation quantifies high-ability and coethnic peer effects in higher education located in an ethnically diverse setting. While empirical research has documented the negative impact ethnic diversity has on several political and economic outcomes in Sub-Saharan Africa, including economic growth, political engagement, conflict, and contributions to public goods, we know relatively little about educational peer effects in such settings, which are generally characterized by high ethnic diversity and cross-ethnic mixing. This chapter studies the effect of coethnic and high-ability peers in student groups on academic outcomes at a large public university in Uganda, a country with pronounced ethnic heterogeneity and segregation. I link data on student-level university admissions with subsequent grades. Upon admission, dorm assignments are random conditional on gender, providing exogenous variation in peer group formation. On average, high-ability peers (irrespective of ethnicity) and coethnic peers (irrespective of ability) positively affect a student's performance. Whereas the coethnic peer effect disappears by the year of graduation, the high-ability peer effect persists and even increases in magnitude over time. Lastly, I find that the effect of high-ability coethnic peers on performance is statistically indistinguishable from that of high-ability noncoethnic peers.

The second essay uses a causal forest algorithm to analyze heterogeneity in coethnic peer effects by estimating a grade dose-response function and treatment effects resulting from interethnic

relations. Specifically, I train the causal model to optimize heterogeneity in students' characteristics, including ethnic groups. This model predicts each student's conditional average treatment of coethnic share while doing data-driven sample splits to estimate heterogeneity. I find that coethnic peer effects are strongest for the largest ethnic group. This is the group that portrays more ethnic attachment than other ethnic groups in this setting. I also find the lowest effects for the second-largest group, which controls the central government and is thus likely to identify more with national identity than tribal identity.

The last essay uses a field experiment to examine how employers select employees for training and the demand for training from employees. Along with collaborators, I elicit employers' beliefs about which of their employees it would be socially optimal to train and their preferences over which employees they choose to train. I then investigate whether employers' selection of workers is individually rational. Finally, I measure employees' self-selection into training and their alignment with employers' selections. To ensure incentive compatibility of employer and employee choices, I provide employees from a sample of metalworking SMEs with free, high-quality skills training. Additionally, I conduct practical skills tests to measure employee metalworking skills before and after training. My analysis shows that owners perceive that training improves the quality of a trained worker. Yet, when offered the opportunity to choose an employee for training, they do not select workers whose quality would improve the most from training. Instead, they choose workers with ties to the firm, as those workers are perceived to be most profitable post-training but would not gain the most from training.

Acknowledgements

This achievement would not have been possible without the support of several individuals, during my intellectual and personal growth. First, I want to thank my committee, Travis Lybbert (chair), Scott Carroll, and Stephen Boucher. They have shown belief in me as a researcher and fostered an environment that has led to tremendous growth during my time at UC Davis. Additionally, Travis and Steve have provided a 'community' outside of school, helping me miss my family a lot less while at Davis, especially during the holidays. I have learned an incredible amount from each of my committee members, both professionally and personally. I am also grateful for their support during the job market process.

Outside my committee, I have obtained support from several Economics and ARE professors during my job market and thesis journey at UC Davis, such as Arman Razae, Dalia Ghanem, and Rachael Goodhue, as well as from development group folks, especially Ashish Shenoy. I also want to acknowledge my co-authors, Andy Brownback, Arman Razae, and Sarojini Hirshleifer. I have learned an incredible amount from them, especially how to apply economic theory in development economics and how to implement and manage field experiments.

My time at UC Davis would not have been possible without a supportive student community. I especially want to thank Scott Sommerville for his friendship; Saloni Chopra for her research collaboration; Jeffrey Hadacheck and James Keeler for their friendship and constant help with coding in R and editing my earlier manuscripts. I did not have any R coding experience at the start of my PhD, but James and Jeff helped me hone my skills. I also want to thank Kavin Dihn for helping me prepare for my micro preliminary exam and Alex Machinda for collaborating and helping me with data scraping.

Lastly, I want to thank my family who have consistently reminded me of the meaning of life and helped me feel loved. First, I want to thank my mom and dad, who, despite only completing Grade Four and barely speaking English, appreciated education's role in breaking the cycle of poverty that has plagued my ancestral home. They instilled in me a strong foundation growing up in southwestern Uganda, where less than one percent attain post-secondary education. My older siblings, Sarah and Henry, took care of me and my younger siblings while my parents were working and taught us the value of hard work. My sister was the first to come to the US for graduate school,

inspiring me that obtaining a PhD from a first-class university was possible, and encouraging me to apply. To my younger brother Peter, for his constant friendship and generous heart, and to my brother-in-law Joe, for his financial support when I first came to the US for graduate studies, and to my other siblings also, thank you.

Contents

List of Tables

List of Figures

Chapter 1

Introduction to the Essays

Human capital accumulation, including education and training, is essential for economic growth. Becker (1964) highlights that investments in education, training, and health boost individual productivity and lifetime earnings. Vocational training programs, supported by various governments, significantly enhance worker skills (King and Palmer 2010). Peer effects also impact human capital development, especially in educational settings. During early education, exogenous factors such as gender (Gong, Lu, and Song 2019) and home environment, such as domestic violence (Carrell and Hoekstra 2010), can influence grades in elementary school. In higher education, students may be more likely to sort themselves despite exogenous peer groups (Carrell, Sacerdote, and West 2013). While research often focuses on Western colleges, where race (Black or White) is a regression control, high ethnic diversity and segregation in Sub-Saharan Africa may complicate and add new dimensions of peer effects.

Uganda is characterized by high ethnic heterogeneity and segregation (Uganda Bureau of Statistics 2016; Alesina and Zhuravskaya 2011). I use a setting at Makerere University Kampala Uganda (MUK), which is centrally located and one of the most prestigious universities in Uganda, to study peer effects and ethnicity in chapters One and Two. By its virtue of location and prestige, MUK enrolls students from disparate ethnic regions of the country. It is, thus, at the MUK campus that most students interact with peers of different ethnicities. Additionally, dorm assignments are random upon admission and within each cohort, providing exogenous assignments to peer groups. Specifically, I define a peer group as students admitted to the same school, such as the School of Medicine, and randomly assigned to the same dorm. My identifying assumption is that, conditional on gender, school, and cohort, a student's peer quality measures (share of coethnic peers and share of high-ability peers) are uncorrelated with unobservables and student characteristics. Specifically, this is the variation that I exploit. Consider two students admitted to majors in the same school and year but randomly assigned to different dorms. One student may be exposed to x coethnic peers in their dorm, and the other student may be exposed to y coethnic peers in another dorm as a result of randomization, where $x \neq y$.

My analysis in Essay One shows that both coethnic peers and high-ability peers are essential for academic performance when students first arrive on campus. Initially, the presence of coethnic peers significantly boosts performance, but this effect disappears by the time students graduate. In contrast, the impact of high-ability peers persists and even increases in magnitude over time. Additionally, I find that the entire effect of peer influence is driven by students with limited prior exposure to different ethnic groups. These are students who have graduated high school in their district of origin and whose ethnic identity is likely more salient as a result of migrating to the city for university in a different ethnic region (Okunogbe 2024). My findings might be explained through peer-to-peer learning channels and psychological channels, such as the contact hypothesis in William (1947). This setting makes it easy for students to learn which of their peers are high-ability and seek study help from them if needed. Also, student-teacher interactions are limited because universities in this setting do not offer office hours, making peer-to-peer learning essential for student success. Lastly, the results showing that coethnic peer effects disappear by graduation imply that as students progress, they navigate diverse environments, experience cross-ethnic interactions, and form cross-ethnic networks, making the psychological boost of coethnic peers less important.

Essay Two builds on the first essay and studies heterogeneity related to coethnic peers. Specifically, I use causal forest estimation methods developed in Athey and Imbens (2016) and Athey and Wager (2021) to analyze the nonlinearities and potential interethnic composition effects. Causal forest uses data-driven sample splits, reducing researcher bias in selecting the relevant heterogeneity dimensions. Additionally, the causal forest method enables the capture of high-dimensional nonlinearities while avoiding overfitting by employing training sampling differently from an estimation sample. Essentially, I estimate Conditional Average Treatment Effects (CATE) for each individual by feeding the causal forest algorithm an estimation regression similar to the primary estimation regression of Chapter One.

I find substantial differences in the estimated CATE by ethnicity. For example, the distribution of the largest ethnic group, Baganda, stochastically dominates the distribution of other ethnicities and the average CATE of this group is larger than the average treatment effect (ATE). In contrast, the distribution of the second largest group, Banyankore, is centered to the left of the ATE. I also find that heterogeneity by gender within each ethnicity is not unidirectional. For example, coethnic peer effects seem to matter more for women in the Basoga, the third largest ethnic group, whereas, for Baganda, coethnic peer effects seem to matter more for men. The patterns of results of this heterogeneity suggest that homophily (as a result of ethnic attachment), not inter-ethnic relations, may be driving coethnic peer effects in this setting.

The last essay studies human capital development in terms of employee training. Motivated by Becker (1962) who predicts an under-provision of general skills training in perfectly competitive markets because owners would have to pay a worker the post-training marginal product. Otherwise, a firm faces the friction of separation by the employee. Because of such market failures, governments in developed countries step in to provide subsidized training. However, the consequences of training under private provision might be more extreme in the developing world, as such settings often suffer from low productivity and imperfectly competitive markets (Hsieh and Klenow 2009).

Together with collaborators, we provide a free training program to employees of small (4-14 workers) metalworking firms in Uganda and study how owners select workers for training. Additionally, we elicit from owners about the perceived gain from training each of their workers and offer objective measures of quality. Since our training is free and carefully designed to limit non-monetary costs, the owners in this evaluation sample should be able to pay a worker their post-training marginal product if they anticipate any separation by the worker after training and should select a worker who would improve the most from training.

Although I find that owners believe our training would improve workers' skills or human capital through our baseline questions and incentive-compatible elicitation for training selection, our data show that owners do not select the workers who would improve most from training. Instead, they select a worker with the strongest ties to the firm. My analysis suggests that owners operate rationally and try to maximize the gap between the marginal product and the wage they can pay their workers without causing that worker to leave. In doing so, they select the worker who is least likely to separate from the firm, often a relative, rather than the worker who would gain the most from training. Lastly, the analysis shows that owners' beliefs do not align with the beliefs of the workers.

References

Alesina, A., and E. Zhuravskaya. 2011. "Segregation and the Quality of Government in a Cross Section of Countries." American Economic Review 101:1872–1911.

Athey, S., and G. Imbens. 2016. "Recursive partitioning for heterogeneous causal effects." Proceedings of the National Academy of Sciences 113:7353–7360.

Athey, S., and S. Wager. 2021. "Policy Learning With Observational Data." Econometrica 89:133–161.

Becker, G.S. 1964. Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education. Chicago: University of Chicago Press.

—. 1962. "Investment in Human Capital: A Theoretical Analysis." Journal of Political Economy 70:9–49.

Carrell, S.E., and M.L. Hoekstra. 2010. "Externalities in the Classroom: How Children Exposed to Domestic Violence Affect Everyone's Kids." American Economic Journal: Applied Economics 2:211–28.

Carrell, S.E., B.I. Sacerdote, and J.E. West. 2013. "From Natural Variation to Optimal Policy? The Importance of Endogenous Peer Group Formation." Econometrica 81:855–882.

Gong, J., Y. Lu, and H. Song. 2019. "Gender Peer Effects on Students' Academic and Noncognitive Outcomes: Evidence and Mechanisms." Journal of Human Resources, pp. .

Hsieh, C.T., and P.J. Klenow. 2009. "Misallocation and Manufacturing TFP in China and India*." The Quarterly Journal of Economics 124:1403–1448.

King, K., and R. Palmer. 2010. Planning for Technical and Vocational Skills Development. Paris: UNESCO.

Okunogbe, O. 2024. "Does Exposure to Other Ethnic Regions Promote National Integration? Evidence from Nigeria." American Economic Journal: Applied Economics 16:157–92.

Uganda Bureau of Statistics. 2016. "The National Population and Housing Census 2014 – Main Report." Working paper, Kampala, Uganda.

William, J., Robin M. 1947. The Reduction of Intergroup Tensions: A Survey of Research on Problems of Ethnic, Racial, and Religious Group Relations. New York: Social Science Research Council.

Chapter 2

Peer Effects and Ethnicity in Uganda

2.1 Introduction

Understanding the determinants of academic and other outcomes for students in higher education continues to be a priority for university administrators and policy-makers. While significant progress has been made in understanding the role of peer effects on academic performance (Sacerdote 2011; Foster 2006; Zimmerman 2003) and other outcomes, such as major choice De Giorgi, Pellizzari, and Redaelli (2010), and cheating (Carrell, Malmstrom, and West 2008). Most of this research has been conducted in the West. Whether or not these results translate to developing countries, such as those of Sub-Saharan Africa (SSA), is unclear. Indeed, since peer effects reflect social dynamics that can change dramatically across cultural contexts, it seems likely that these effects could operate quite differently in non-Western settings. One specific reason to doubt the external validity of the existing peer effects literature on SSA is the degree and nature of ethnic diversity that characterizes much of the region. Such heterogeneity combined with ethno-linguistic differences may, for example, complicate student collaboration, thereby muting the positive effects of high-ability peers on student performance.

Uganda, the setting for this study, consists of over 50 ethnicities (Uganda Bureau of Statistics 2016). These ethnicities are geographically segregated, although there is considerable ethnic mixing in the capital of Kampala. Several studies link high ethnic heterogeneity in SSA to several poor economic outcomes, such as public goods provision, economic growth, and firm productivity, and to negative effects on social indicators, especially social trust.[1] A prevalent bias in favor of coethnic interaction partly explains these documented costs associated with high ethnic diversity in SSA. Coupling this diversity with strong ethnic segregation, as is the case in Uganda, further exacerbates mistrust (Alesina and Zhuravskaya 2011). This added friction to social interaction, cooperation, and collaboration is costly in general but may be especially apparent in student performance at universities that draw from disparate ethnic regions and, hence, where many students first interact intensively with ethnicities other than their own.

This paper leverages the higher education context in ethnically diverse and segregated Uganda to explore the effects of coethnic and high-ability peers on academic outcomes. This unique empirical setting raises a number of questions that this paper studies. Does the share of coethnic peers within a student's peer group affect academic performance more or less than the share of high-ability peers? Do high-ability coethnic peers matter more than high-ability noncoethnic peers? Does the context of Ugandan higher education translate into coethnic peer effects stronger for some students than others? The contribution of this paper to the peer effects literature hinges on providing credible answers to these questions in this distinctive setting.

In the empirical stage for this analysis, I link administrative records of student applications, admissions, and post-admission academic performance from a large public university in Uganda. These records include students enrolled in most of the STEM, social sciences, and business degrees in the years 2009-2017 at this prestigious national university that is centrally located and, by admitting students from across the country, creates a microcosm of Uganda's rich ethnic heterogeneity. For the purposes of this research, this feature is particularly interesting given the strong geographic segregation of ethnicities in Uganda, which means that many students arrive at the university with little prior exposure to other ethnicities but are suddenly surrounded by the full diversity that constitutes the country as a whole. In the analysis that follows, I classify students who graduated high school from their districts of origin as those with less prior exposure to other ethnicities and for whom the ethnic diversity on campus is most salient.

Being surrounded by coethnic peers at this large university might provide a sense of belonging

[1] For example, cross country quality of government (Alesina and Zhuravskaya 2011); cross country public policies (Easterly and Levine 1997); productivity of a firm in Kenya (Hjort 2014); public goods provision in Uganda (Habyarimana et al. 2007) Additionally, regarding public goods, Gisselquist, Leiderer, and Niño-Zarazúa (2016) show that high ethnic diversity may lead to welfare gains. For ethnicity and social trust, see Alesina and La Ferrara (2000)

and stability, thereby enhancing academic performance. Interacting with high-ability students can similarly improve performance in this setting because contact with instructors is limited (e.g., office hours are not offered), so learning from peers is important. In addition to testing the direct effects of coethnic and high-ability peers, I also estimate the interaction effect of these two peer types since homophilous coethnic sorting could hamper or help learning from peers depending on the academic ability of these coethnic peers. In this analysis, I rely on exogenous variation in the share of coethnic and high-ability peers in a given student's peer group to test for these direct and interaction effects.

The administrative data I use in this paper provide students' demographic and academic characteristics, including whether they were admitted on merit scholarships, which I take as an indicator of high ability. These records do not, however, report student ethnicity. I overcome this limitation by exploiting linguistic and cultural characteristics common to Uganda and SSA, where surnames reflect one's native languages and, thus, ethnicity. To do so, I apply a machine learning algorithm common in computational linguistics introduced in Cavnar and Trenkle (1994) and recently adapted by **?** to the Ugandan context to a national administrative dataset of 2016 voter registrations that includes over 14 million Ugandans. This external data set provides training data I use to build a classification model that predicts ethnicities using student surnames.

This paper's causal identification of peer effects hinges on the random assignment of incoming students into dorms, which provides exogenous variation in peer groups. Specifically, a peer group in this analysis consists of students admitted to majors in the same school and assigned to live in the same dorm. Upon admission and conditional on gender, dorm assignment is random. Since there is excess demand for dorm beds, actual residence in dorms is not guaranteed. Some end up living off-campus, but dorm assignments shape campus life for some of these students, as they may engage in extracurricular activities within their assigned dorm. In addition to exogenous assignment to peer groups, the econometric strategy exploits idiosyncratic year-to-year variation in coethnic composition. Moreover, each student's course list is predetermined at the time of admission, and students do not meet their classmates and dormmates until orientation week. Thus, the results in this paper are not driven by selection into peer groups. I control for dorm, classroom (course-by-year), and major fixed effects to account for correlated shocks and differences that might confound my estimates.

The results of this analysis indicate that the coethnic peers are as important as high-ability peers in this setting, especially in the first year. That is, I find that coethnic peers (irrespective of ability) and high-ability peers (irrespective of ethnicity) increase a student's performance in the first year. Specifically, adding five coethnic peers to a peer group of size 25, which would increase the number of coethnic peers in a group from the twenty-fifth to the seventy-fifth percentile, increases a student's performance by 0.19 percentage points. Additionally, the same change of adding high-ability peers to a group of 25 increases a student's performance by 0.15 percentage points. Both effects are significant at the 5% level and are about 0.02 standard deviation change in a student's performance in the first year.

Nevertheless, the effect of coethnic peers disappears by the time a student graduates but that of high-ability peers persists and even increases. Specifically, the effect of coethnic peers in the third year, which is the final year for almost all the majors in this setting, is half of that observed in the first year. Yet the effect of high-ability peers in the third year is 1.5 times that of the first year in magnitude. Lastly, although I find that suggestive evidence shows that coethnic peers matter more than high-ability noncoethnic peers as a student advances during university education, the effects of both types of peers are statistically indistinguishable.

Beyond the average effects, heterogeneous impacts indicate that the effect of coethnic share is mostly driven by students of assumed high ethnic salience. For example, adding five coethnic peers to a group of 25 increases the academic performance of a high ethnic salience student by 0.05 standard deviations, which is 2.5 times the average effect. Nevertheless, like the mean effect of coethnic peers, this effect on students with high ethnic salience fades as a student progresses. On the contrary, coethnic peers have a positive and significant effect on high-ability, not low-ability, students that persist into the third year, suggesting that the benefits of coethnic peers throughout a student's university career can be reaped by those posed to succeed when they enter university.

Qualitative insights from the specific university context of this study align with potential underlying explanations for these results, including peer-to-peer learning and cultural and psychological factors. In this setting, both coethnicity and academic ability are readily and generally observable. Incoming freshmen can easily identify coethnic peers through physical features and cultural characteristics, including names and language. As the academic year unfolds, they also learn who among their peers are high-ability because publicly posted scores and grades reveal academic merit scholarship status or through frequent interactions. Given the prevalence of ethnic student organizations and activities on campus, which suggests a degree of homophily that shapes student life, it is natural for incoming students to seek out coethnic connections and support. Such connections can be critical to a student's successful transition to a novel setting of high ethnic diversity and, possibly, latent inter-ethnic tensions that may prevail on campus.

This university setting is also characterized by classical lecture-style instruction with few opportunities to interact with faculty or consult with teaching assistants, which makes informal peer-to-peer learning especially important. For both incoming and continuing students having high-ability peers in one's peer group can therefore provide an advantage. If anything, the benefit of such informal peer tutoring increases as students progress to more advanced courses in their degree programs.

The finding that a higher coethnic share matters on average and especially so for students of higher ethnic salience is suggestive of other psychological mechanisms. Enrolling at a large, centrally located national university may increase ethnic identity salience and attachment, as social identity theory in Tajfel (1982) predicts. The presence of coethnic peers in one's university environment may thus be beneficial for such students. This is similar to the finding reported in Okunogbe (2018), showing that the ethnic pride of Nigerian youth increases when they do national service in a region where they are not part of the ethnic majority. Also, since several students are forced to navigate a space that is diverse compared to their pre-university schools, having coethnic peers precludes inter-ethnic barriers. Moreover, I find the differential effect of the share of coethnic peers on students of high ethnic salience observed in the first year disappears as a student progresses. This indicates that through frequent cross-ethnic interactions at the university, these types of students make cross-ethnic networks and factors other than the shared ethnic identity of peers begin to matter more for academic performance. This phenomenon can be interpreted by the contact hypothesis in William (1947). This might also explain why the effect of high-ability peers increases over time.

This paper contributes to several strands of literature on peer effects in college. Although mixed, prior evidence largely indicates that post-secondary peer effects meaningfully impact education outcomes, such as major choice and academic performance. For example, Zimmerman (2003) and Sacerdote (2001) exploit random roommate assignments at US colleges to study roommate peer effects. Zimmerman (2003) finds significant but small peer effects when using pre-treatment academic characteristics to measure peer quality and also detects nonlinear effects that are conditional on the student's SAT scores. Sacerdote (2001) finds null effects using the ability of a peer but significant nonlinear effects at Dartmouth on academic outcomes. Additionally, he finds strong effects on some social outcomes (e.g., fraternity membership). Carrell, Fullerton, and West (2009) argues that roommates are a small part of one's college life, which might explain why Zimmerman (2003) and Sacerdote (2001) find no strong dorm or roommate peer effects. Exploiting exogenous assignments at the United States Air Force Academy, where students are assigned to peers with whom they spend a majority of time together, Carrell, Fullerton, and West (2009) find stronger academic peer effects than roommate peer effects. More recently, Mehta, Stinebrickner, and Stinebrickner (2018) use a panel data set that tracks students' time allocation and friendships at Berea College and found that peers have an effect on study efforts.

These peer effects studies primarily focus on college peers in the West. Their main econometric specifications include the average quality of peers measured by pre-treatment academic characteristics on the right-hand side variables. Given the setting, these papers also control for race, usually a binary indicator for white or black. In the SSA region, however, high ethnic diversity introduces new complexity and nuance to peer effects. For high-ability noncoethnic might have a negative or null effect on academic performance if high ethnic diversity leads to inter-ethnic rivalries and discrimination that spill into classrooms. I find the opposite: the identity of a high-ability peer does not matter. High-ability peers (irrespective of ethnicity) affect a student's academic performance, suggesting that peer effects observed in studies in the West also exist in this setting. I find that coethnic peers are also important in the first and second years.

This paper also contributes to the literature exploring the role of ethnic diversity on economic and social outcomes in SSA more broadly (Easterly and Levine 1997; Habyarimana et al. 2007; Alesina and Zhuravskaya 2011; Miguel 2004; Gisselquist, Leiderer, and Niño-Zarazúa 2016; Alesina and La Ferrara 2000; Håkansson and Sjöholm 2007; Hooghe 2007). In contrast to these more general studies, this analysis focuses on a different question, albeit one with clear importance and policy relevance. Understanding peer effects from social networks play out in higher education institutions with high ethnic diversity may enable more informed admissions and other academic processes, which often feature explicitly or implicitly in facilitating (or potentially undermining) cross-ethnic cooperation among young adults. High-ability peers affecting academic outcomes more than coethnic peers as students progress may suggest that Ugandan youth are less ethnically biased or able to adapt to ethnic diversity. However, it is important to note that coethnic peers might have lasting impacts on social networks outside school or other outcomes that are unavailable in my data.

Peer effects in higher education in SSA have been understudied for several reasons, including data constraints. A few studies that have explored college peer effects in the region use data from a South African university (Garlick 2018; Corno, Ferrara, and Burns 2019). Nevertheless, Garlick (2018) focuses on peer effects under two different assignment rules (random and residential tracking), while Corno, Ferrara, and Burns (2019) focuses on how exposure to roommates of another race changes one's stereotypes. Race (white vs black) is salient in South Africa for historical reasons and general population composition, unlike other African countries. Therefore, I add to the literature by studying higher education peer effects at a university in a context about which we know very little. I find that in this setting, coethnic and high-ability peer effects exist, especially in the first year.

2.2 Background

2.2.1 Ethnicity in Uganda

Uganda has over 50 ethnic groups that belong to three broader Bantu-speaking tribes (UBOS 2006). The largest nine ethnic groups constituted 71% of the population according to the 2002 Uganda Population and Housing Census. Groups may differ by traditions (e.g., dressing), language, food, economic activities, and sometimes by physical characteristics (e.g., skin tone). This pronounced ethnic diversity is also characterized by distinct geographic segregation as shown in Figure 2.1. Indeed, Alesina and Zhuravskaya (2011) rank Uganda the 4^{th} most segregated countries in the world based on a spatial segregation index.

Historic migration and ethnic kingdoms drive these segregated settlement patterns. Bantu-speaking groups are clustered in the country's South, Central, and Western parts, while Nilotic and Nilo Hamites peoples are clustered in the Northern and Eastern parts. For purposes of the analysis that follows, I retrace current ethnic borders to historic kingdoms (see Appendix Section 2.8.2). Inter-region migration is limited except for rural-to-urban migration into the capital, Kampala, for economic opportunities. By contrast, rural-to-rural migration across ethnic clusters rarely opens

Figure 2.1: Geographic Segregation.

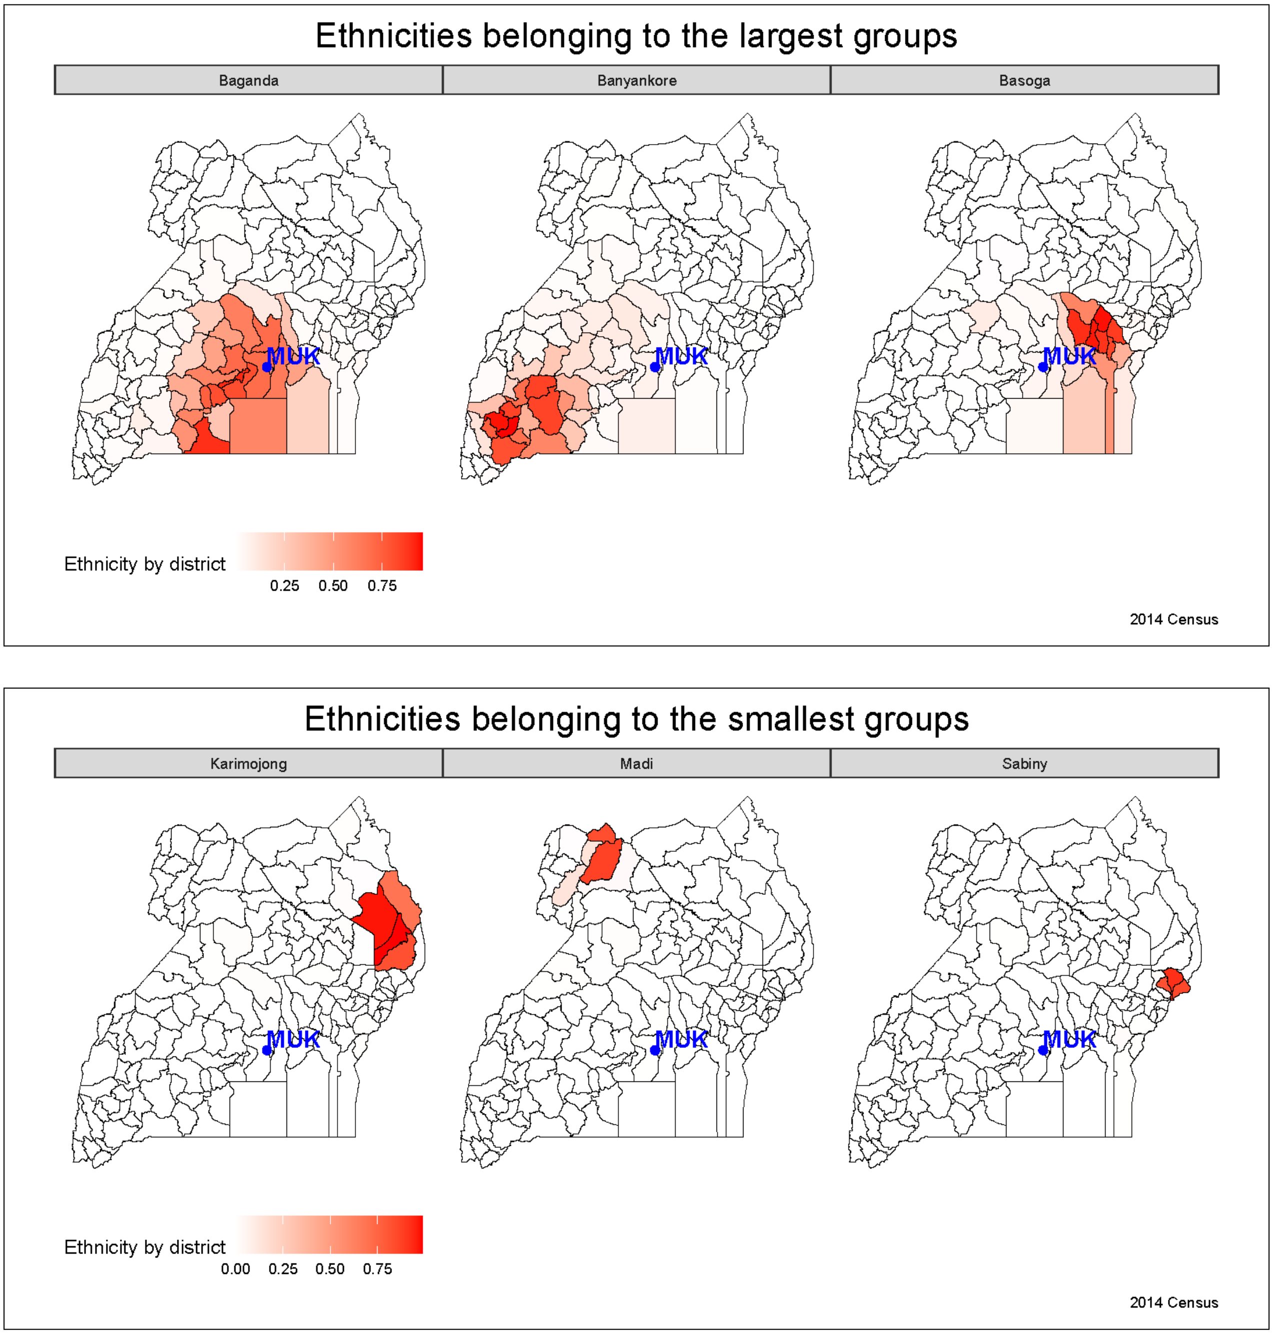

Notes: Ethnicity by district is the proportion of each ethnicity within a district. Data source: 2014 census. District shape files can be downloaded from https://data2.unhcr.org/en/documents/details/83043.

economic opportunities and is limited due to cultural reasons.

Although ethnic divisions existed in pre-colonial Uganda, some were exacerbated during British colonialism (Tornberg 2013). The first post-independence government made efforts to reduce the importance of ethnic identities by abolishing historic kingdoms and preaching national unity, an effort that met with resistance from some kingdoms, especially those with economic or political power. The current government allowed ethnic groups to reinstate their historical kingdom; some ethnic groups did. While current inter-ethnic competition and recent historical conflicts can be traced to political and sometimes historical factors (Mamdani 2001), inter-ethnic competition or outright conflict is generally not as intense as in neighboring countries.

Although English, the official language of Uganda, is spoken in public offices and taught in schools, native linguistic diversity is high.[2] Differences between native languages are correlated with physical distance, implying that one may partially comprehend the language of a neighboring tribe. Luganda is the most familiar native language because it is native to the Kampala region. I exploit this language diversity to predict ethnicity in Section 4.

2.2.2 Ugandan Higher Education and Makerere University Kampala

Although Uganda has one of the youngest populations in the world, post-secondary school education is low: the post-secondary enrollment rate for college-age Ugandans was only 6.85% during the 2017/18 academic year (NCHE 2018). Nine public and 44 private universities offered degree programs during the 2018/19 academic year (NCHE 2018), of which Makerere University Kampala (MUK) ranks first in quality and size.

MUK is well-known in the SSA as it is one of the oldest universities in the region. It was established in 1922 as a technical school to facilitate training workers for the British colonial government. It is centrally located in Kampala and admits students from across the country. For some students, it is at this university that they meet and interact with people of different ethnicities for the first time. With the exception of Baganda, the diversity of the MUK student population mirrors that of the country as a whole (see Figure 2.2).

[2] WorldAtlas reports Uganda's language diversity index of 0.929, which indicates that most Ugandans speak at least one native language.

Figure 2.2: Distribution by ethnicity: student sample vs general population

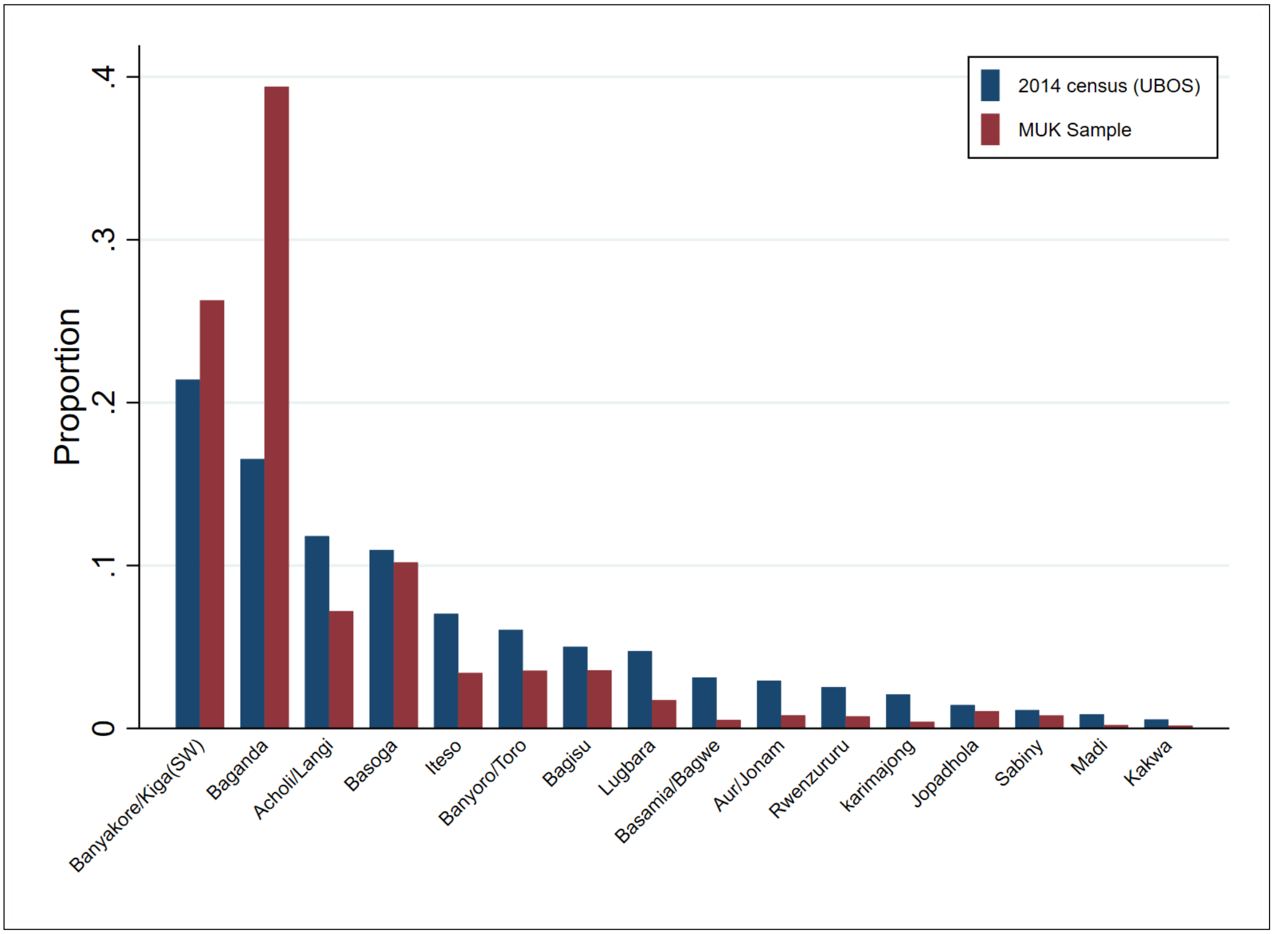

Source: MUK admissions, 2009-2017 and 2014 Census (UBOS)

2.3 Empirical Setting

2.3.1 Applications, Admissions, and Sample Definition

The Ugandan public university and pre-collegiate nationwide system offer a unique setting that I exploit to identify coethnic peer effects. First, national pre-collegiate exit exams and public university merit scholarships are centrally administered. Second, the Uganda Examination Board, an organization separate from MUK, runs an algorithm for all MUK admissions. Thus, there is no room to manipulate the composition of its student population.

Students are admitted under two schemes: (I) National merit scholarship and (II) self-funding scheme. A student lists up to six majors in order of preference during application. Admission to a major (cutoffs) is a function of the student's preference set, admission in national exams, and the university's capacity. A student's major (and course list) are predetermined during admission, 3-4 months before enrolling. Each major non-extension major is housed within a school, which is a smaller unit within a college. A school is locally termed as faculty or department, but I will adopt the 'school' term for simplicity.

Students in the same majors take almost all their first year classes together since courses have predetermined sequencing. Still, they interact with students from other majors within classrooms, usually within the same school, who share the same course requirements on a daily basis. In addition, students within a school usually share common spaces, such as computer labs, food canteens, study rooms, and libraries.

Students cannot select into different sections within the same major, as sections do not exist in this setting. Because of this, most student's course sequence is also pre-determined before a student reports to campus. The performance data show that over 98% of classes in each year are non-elective. Moreover, The university offers evening and day class options as 'different' majors when a degree, such as business administration, is in demand. Still, the day class is a 'different' major from the evening class, and students must apply and get admitted to either the day or evening class cohort separately. For example, students who intend to obtain a Bachelor of Business Administration degree can apply and be admitted to either the day cohort or the evening cohort. Students admitted to the day cohort cannot take classes and sit for their exams with students admitted to the evening cohort. I restrict the sample to day cohorts as evening majors do not qualify for the national merit scholarship. This is important because the merit scholarship is my measure of high-ability as I define in the coming sections.

Students stick with the majors offered during admissions but can apply to change within the first two weeks of their freshman year. Approvals depend on the capacity of the intended major and student performance and are thus rare. I find major change cases are less than 2.75% in the ten-year period of my sample.[3] Non-STEM majors, especially business and social sciences, tend to have relatively large class sizes.

[3]I compare the major student's enrollment and the major at the admissions and find a mismatch of 2.75%. This number includes students whose major switch applications were approved and possibly some data entry errors when entering admissions data.

2.3.2 Dorm Assignment and Defining Peer Groups

Conditional on gender and upon admission, dorm assignment is random. MUK has nine single-sex large dorms: three are female and six are male dorms. There are more incoming students assigned to dorms than there are beds to accommodate them. I observe dorm assignments but not the subsequent residence status and room assignments. Each student's admission letter indicates the assigned hall, which determined by the administration by simple random assignment. Students must formally apply to their assigned dorm for residence, at which point a dorm administrator and committee allocates beds according to a university-wide priority list that favors students on national merit scholarships in majors and schools perceived to be especially rigorous, such as medicine and engineering.

The remaining beds are then assigned to students according to the order of their dorm application. While students who are not allocated a bed in their assigned dorm must arrange for their own housing off-campus, their initial dorm assignment continues to shape campus life as assigned students have access to shared spaces with entertainment and dining facilities in these dorms. Extracurricular activities such as student government elections are also organized by dorm assignment irrespective of residence.

In general, a peer group consists of individuals with shared or similar characteristics who interact in social or other settings. Specific definitions of peer groups are context-specific. Carrell, Fullerton, and West (2009) define a peer group as a squadron at the US Air Force Academy, while Foster (2006) consider a peer group to be students living on the same dorm floor and Pre-collegiate studies, such as Carrell and Hoekstra (2010), use a cohort definition. In the MUK context, I define a peer group, as illustrated in Figure 2.3, as students within the same school who are assigned to the same dorm. Although most of the prior studies on college peer effects observe roommates, I do not observe room allocations and residence status, so I restrict my definition to dorm assignment.

By focusing on the cohort-residential peer groups, I use a "strict" definition of a peer group, but it also allows me to study peers with whom a student spends most of their time. For instance, students within the same school and dorm may spend a lot of time together, such as walking to and from classes, attending classes together, and sharing common spaces at school and within the dorm. One may argue that a major hall year is a better peer group since students in the same major take

Figure 2.3: Peer Group Construction

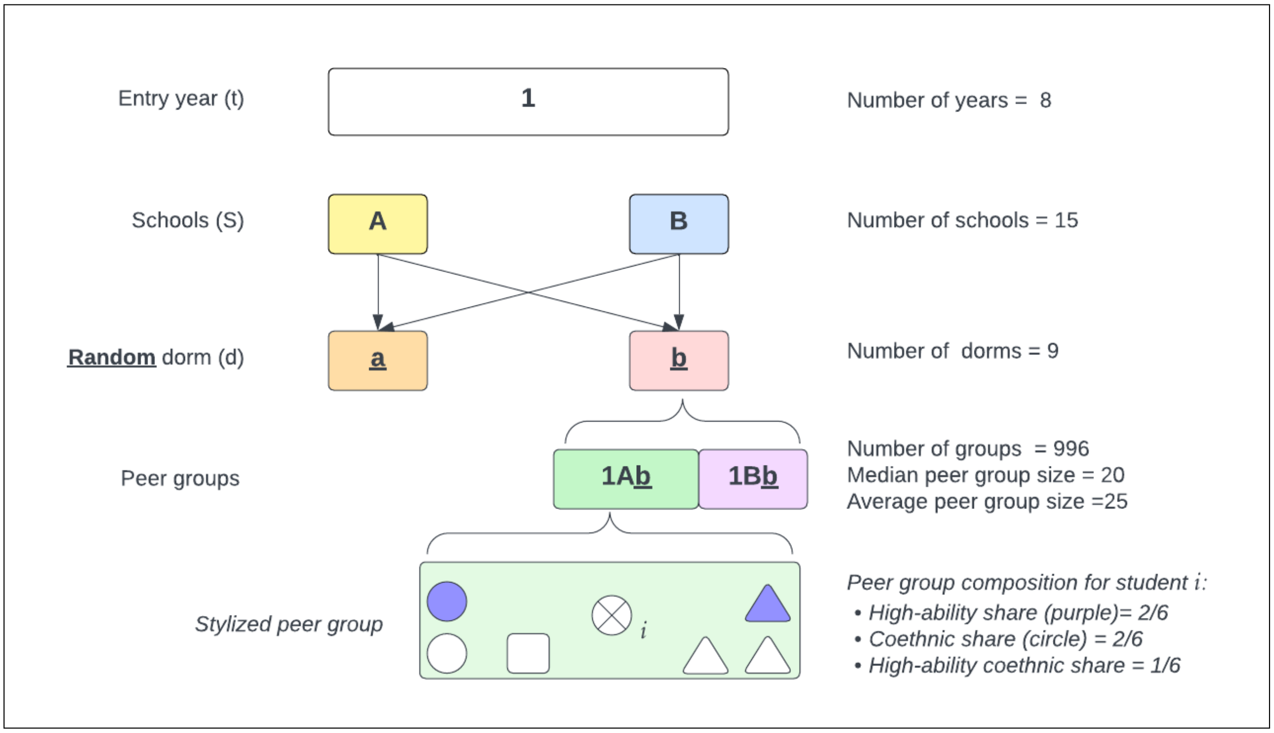

Notes: This diagram illustrates a peer group definition. Students in these defined peer groups are much more likely to interact regularly with each other, including those of the same or different ethnicity. High-ability students are defined as those on merit scholarships, a status that is widely known among all students.

100% of their classes. However, the focus of this paper is coethnic and high-ability shares, and using a way smaller peer group definition reduces variation in the coethnic share as most of the coethnic share of the smallest ethnicities will be zero in a lot of peer groups.

2.3.3 Identifying Peer Effects at MUK

Estimating peer effects may be econometrically challenging for three reasons: self-selection (Hoxby 2002), endogeneity (Manski 1993), and correlated common shocks (Bramoullé, Djebbari, and Fortin 2009). This section highlights the characteristics of this setting that provide solutions to these issues related to measuring peer effects.

Self-selection arises when people choose to join a group based on some pre-treatment characteristics. As stated in Hoxby (2002) *".. if everyone in a group is high achieving, many observers assume that achievement is an effect of belonging to the group instead of a reason for belonging to it."* In the case of colleges, self-selection exists because students can select into classrooms, majors, and sometimes, dormitories. Peer effects literature typically employs two strategies to deal with selection in peer effects papers. First, conditional on some pre-treatment characteristics, such as gender and ability, peers result from random assignment (Sacerdote 2001; Zimmerman 2003; Carrell, Fullerton, and West 2009; Foster 2006). However, random assignments into classrooms in US studies are difficult. So, with the exception of Carrell, Fullerton, and West (2009), these studies use a setting where roommates at some universities are randomly assigned.

The second approach involves exploiting natural variation in a cohort or group composition. The idea behind this approach is that year-to-year variation (e.g., gender, race, and class size) observed at the group level is a reflection of a natural variation in a general population (idiosyncratic). This approach has been used in pre-collegiate peer effects studies (Carrell, Hoekstra, and Kuka 2018; Hoxby 2000)

My approach leverages characteristics of this setting described in Section 2.3.1. Peers are classmates who potentially live together. As aforementioned, conditional gender dorm assignment at MUK is random. Since I do not observe roommate assignments and residence status, this paper estimates the intent to treat (ITT) of the peer qualities defined later. Unlike most US universities, students do not select courses or majors post-admission, which has the convenient feature that students do not sort into classes or classrooms based on characteristics or exposure (or not) to different types of peers.

The reflection problem is the endogeneity problem challenge, which arises from a feedback loop of peers. This is a challenge because a student's and their peers' outcomes are simultaneously determined. One of the approaches in the literature is to use preexisting characteristics that are exogenous to the dependent variable, such as race and gender. For example, Carrell and Hoekstra (2010) uses the presence of family problems when studying peer effects of children linked to domestic violence on academic outcomes. I use pre-collegiate characteristics, as most of the literature, to exploit exogenous variation in treatment variables, which are coethnic share and high-ability share within a student's peer group.

In Uganda, students' ethnic identities are determined at birth. An argument may be made that ethnicity is part of multifaceted identities, a function of collective cultural traits, and that an individual's ethnicity may change through self-identification (Sen and Wasow 2016). I am not concerned that this exists in Uganda to the extent that it would confound my estimates. First, I follow the official categorizations of ethnic groups in UBOS (2006), and admissions do not have ethnic quotas or any form of affirmative action based on ethnicity. Thus, there is no incentive to change one's ethnic identity during university applications. Second, I use linguistic characteristics to predict ethnicity instead of self-identification. I describe these variables in Section 2.4.2 below.

The last main challenge is contemporaneous common shocks, especially if they are correlated with academic performance. My setting uses random assignments at the same university, which reduces the possibility of such shocks. Nevertheless, there may be shocks that affect some peer groups differently. Thus, the main regressions include all group fixed effects, such as dorm and classroom, to account for observed characteristics that might confound the main effect.

2.4 The Data

2.4.1 Academic and Demographic Characteristics

Pre-university Characteristics

The analysis in this paper uses several data sources: MUK's administrative records on academic and demographic characteristics observed from applications, admissions, and post-admission academic performance for students entering the university during 2009-2017, and ethnicity is predicted by student surnames.

I observe students' application data from 2009 to 2017. The student applications include the student's name, type of application, admission scheme (merit scholarship or private scheme), and offered majors, as well as age and religious identity. All student records are de-identified pre-analysis, although most student admission data, such as major, are publicized on university notice boards and in newspapers.

Measure of High-ability

Every year, 4,000 students are admitted to public universities on a government merit sponsorship basis of performance in high school national exams, most of which enroll at MUK relative to other public universities (HESFB Uganda, 2012). These scholarships are awarded to the top students within a major, and the number of spots per major is relatively constant across years. Merit scholarship application forms are submitted at the time of national exam registration before students take their exams. Therefore, almost all A-level graduates are automatically considered for the government merit scholarship, as the sponsorship does not require a separate application. Students are ranked based on their high school GPA within their preference set, and the top students are offered a scholarship until each major's scholarship spots are filled up. That is, the scholarship is determined by high school GPA.[4]

High school GPA is a proxy for ability as it may pick up a student's innate ability, effort during high school, and success in an academic context. I therefore use this as an imperfect but informative proxy of "academic potential," which accounts for both a student's subject combination and performance in this selected subject combination. Each major has a high school subject combination required for a student's successful college career from a university's perspective. I define "high-ability" students within each major as those enrolled with the national merit scholarship. Lastly, it is usually public knowledge which of a student's peers are admitted through merit as university registration numbers differ by merit status. Moreover, admission lists are usually published in newspapers and university notice boards.

University Academic performance

I observe student transcripts from 14 departments belonging to six colleges. Each student's transcript lists all courses taken, credit units, and performance in percentages by semester year of study during which the course was taken. Therefore, I can observe these students' classmates and how they have progressed from matriculation to completion of coursework. Unlike schools in the West, letter-grade ranges assignment is the same across all majors, and professors do curve grades. Professors at MUK do not assign letter grades. They submit each student's performance on a 0-100% scale, and the central system assigns the letter grades. Also, most majors take three years to complete, and thus students take a lot of courses per semester (a min of six, and some majors require students to take up to ten courses in some semesters).

[4]There are a few variations. For example, Ugandan public universities have a gender affirmative action policy that awards a 'free' 1.5 additional points to every girl during admission. This 1.5 free point is also awarded to girls when they are being considered for non-merit admission schemes. In addition, a small proportion of the merit scholarship is awarded through the district quota to the top four students graduating from their district of origin who did not obtain the merit scholarship through the direct route. Therefore, the number of district quota spots is proportional to ethnicity size. District quota applications are made at the same time as the national merit applications.

2.4.2 Ethnicity

University applications and admissions do not capture the ethnic identity of students, although ethnicity is one of the most salient identities among Ugandans. I overcome this by exploiting linguistic differences reflected in surnames. Ugandans' surnames are in their native languages.[5] This naming pattern is not random or unique to Uganda. Historically African parents chose names intentionally. However, with the arrival of colonists, first names are now in foreign languages, such as English (in Anglophone countries) or French (Francophone countries). The meanings of most Ugandan surnames can be traced to the father's tribal clan and religiosity or prevailing conditions at the time of birth, among others. These are linguistic characteristics I use to predict one's ethnicity. Data Appendix 2.8.2 describes how I trace ethnic boundaries from current administrative units to historical kingdoms.

Using surnames to trace one's identity is not new in economics and other fields. For example, surnames have been used in mobility studies to trace wealth across generations within a family in the West (Barone and Mocetti 2016; Clark and Cummins 2015). Some studies have also used surnames to predict ethnic identity across several countries. For instance, Bhusal et al. (2020) use surname frequency in the Nepalese historical censuses to predict one's caste in their paper studying how revolutions may have altered political representations and inclusion in Nepal. Using fuzzy matching and naïve Bayes machine learning techniques on historical records, Monasterio (2017) studies surnames and ancestry in Brazil.

Predicting Ethnicity and Constructing Coethnic Share

More related, **?** exploits rural-urban linkages in Uganda and applies machine learning on representative Uganda surnames to predict the rural origin of Uber drivers in Kampala. His study explores how Uber drivers adjust their online hours when their probable ancestral homes experience a negative weather shock. Therefore, agroecological zones form a basis for his predictions. His procedure, like the machine learning section in Monasterio (2017), follows a text categorization procedure developed by Cavnar and Trenkle (1994). This process has been widely used in computational linguistics and involves breaking down a name into N-grams.

[5]Trevor Noah mentions the same pattern in South Africa in his book “Born a Crime" (PP.). Also, see this https://www.bbc.com/news/world-africa-37912748 for another example

Following the literature, there are two common approaches: use frequencies to predict probabilities as in Bhusal et al. (2020) and train a machine learning algorithm on some training data set by applying tools, such as gradient boosting. Method (I) computes simple probability using the frequency of each surname. Suppose $\{E_1, E_2, ..., E_n\}$ is a set of all ethnicities in a population. Also, suppose $N_{s\in E_i}$ is the number of times a surname, s, belongs to an ethnicity, E_i. Then the probability of belonging to a particular ethnicity is computed as:

$$\frac{N_{s\in E_i}}{\sum_{\forall n} N_{s\in E_i}} \tag{2.1}$$

To illustrate, consider the surname "AHIMBISIBWE": it appears 17,559 times in the name training data, of which 13,904 occurrences in the Ankole region/ethnicity. Therefore, there is a 79.2% probability that a student with the surname "AHIMBISIBWE" is of Ankole ethnicity.

Method (II), which is my preferred, follows **?** and computational linguistics and begins by breaking a surname into N-grams. Taking "AHIMBISIBWE" as an example, Method (II) breaks this surname into 1-grams ("A", "H", " M", "B", "I", etc); 2-grams ("AH", "HM", "MB", "BI", etc); 3-grams ("AHI", "HIM", "IMB", etc..) and so forth. The algorithm can now count the number of frequencies each n-gram appears in a surname and in each region. The most common weighting approach used in linguistics is the term frequency-inverse document frequency (tf-idf) that combines approaches developed by Luhn (1957) and Jones (2004). I then apply gradient boosting on N-grams and tf-idf features on an external data set described in Section 2.4.2, producing a classification model that I apply to students' surnames.

The second method is preferred to the first in this paper for two reasons. First, by following the tf-idf weighting procedure, the algorithm picks each surname's most unique linguistic characteristics. Second, it does not require an exact match in the name database.

This algorithm predicts N probabilities if we have N ethnicities. Given that languages are not exclusively unique, some probabilities are non-zero or 1. We can, therefore, interpret the predicted probabilities as a measure of 'confidence' that a student belongs to a particular ethnicity. One approach is to use ethnicity corresponding to the top predicted ethnicity (the ethnicity a prediction is most confident about), as is common in the literature.

I can then compute the share of coethnic peers using two approaches: (A) and (B). Firstly, by single ethnic identity assignment (A), I assign an individual a single ethnicity category corresponding to the group the algorithm is most confident about. This is common in studies employing machine learning algorithms to predict ethnicity, religion, or area of origin in the literature. This method assumes that individuals' top predicted ethnicity corresponds to the 'true' ethnicity and treats ethnicity as a categorical variable without considering potential measurement errors. Using top predicted is common in the literature. The average probability corresponding to ethnicity in the algorithm is most confident equals .792 (median=.861), which is high.

Secondly, by joint probability estimation, I consider all the probabilities that a given surname belongs to different ethnicities. This method acknowledges potential measurement errors associated with using categorical variables for ethnicity. It estimates the probable fraction of coethnic peers in a peer group by considering the joint probabilities of two individuals belonging to the same group.

That is, student i's share of coethnic peers in a peer group G, S^E_{iG} is computed as:

$$\text{Using category assignment (A)}: \quad S^E_{iG} = \frac{\sum_{k\forall\neq i} \text{Number of coethnics}}{N_G - 1} \tag{2.2}$$

$$\text{Using joint probability estimation (B)}: \quad S^E_{iG} = \frac{\sum_{e=1}^{16} \sum_{\forall k\neq i}^{N_G-1} \Pi_{e_i}\Pi_{e'i}}{N_G - 1}, \tag{2.3}$$

where N_G is the peer group size, Π_{ei} is the predicted probability that an individual i belongs to ethnicity group e. Lastly, I collapse ethnicities to 16 ethnic/language groups as described in Appendix 2.8.2. The main analysis uses the probable coethnic share in a group in equation (2.3), but the results remain unchanged when I use the share of coethnic peers computed using equation (2.2). Throughout this paper, I use the share of coethnic peers and the probable share of coethnic peers synonymously for simplicity and ethnicity to mean the most probable ethnicity in the empirical and results sections.

Training Data

I use nationwide voter registration to train the machine learning model (gradient boosted). These data contain names, voter ID numbers, date of birth, sex, polling station, and area of residence. The area of residence is given for all units of administration parish, sub-county, county, and district. I link these voter data to spatial administrative and public data containing ethnic boundaries traced from historic kingdoms described in Appendix 2.8.2.

People register to vote from a polling station within their parish of residence. Moreover, in many cases, people who live in cities outside their areas of origin often register to vote in their ancestral homes. Because voter registration is manual, it only takes place once between elections. A few Ugandans own cars to travel, so most walk, as long-distance public transportation is costly. Thus, the cost of registering to vote in a village different from their residence is high.

2.4.3 Descriptive Statistics

Table 2.1 provides summary statistics for demographic and academic characteristics in Panel A and peer group averages in Panel B. About half of the student population is female, and about 31% are high-ability (enrolled through national merit scholarships). Most of the students have declared religion, and as expected in this context, most students are either Catholic or Anglican. The average age of incoming students is 20. Lastly, about 36% of the students in the sample graduated from a high school with their home district (these are the type of students I assume to have higher ethnic salience). Because of Uganda's high ethnic segregation (Figure 2.1), this group of students may not have interacted with peers of different ethnic groups.

Given my peer group definition, the average group size is 25, which is small, albeit the SD for peer group and cohort sizes are large. Close to 75% of the peer groups are of size 50 and below, as the Appendix Figure 2.6 portrays. STEM majors, which comprise most of my sample majors, usually admit a few students relative to other majors.

The average coethnic share is 25%, and the average co-ethnic share on merit is 7%. To explore how treatment intensity (shared ethnicity) may vary by ethnicity and if MUK is representative of Uganda's ethnicity distribution, I plot the distribution of ethnicities in Figure 2.4. The CDFs in this figure show that 80% of peer groups have a coethnic share of 0.2 or less.

Table 2.1: Descriptive Statistics

	N	Mean	SD
Panel A: Student characteristics			
High-ability	25,487	0.31	0.47
Age	25,487	20.16	1.43
Female	25,487	0.48	0.50
High ethnic salience	25,487	0.36	0.48
Anglican	25,487	0.37	0.47
Catholic	25,487	0.31	0.47
Muslim	25,487	0.09	0.29
Pentecostal	25,487	0.06	0.24
Seventh Adventist	25,487	0.02	0.13
Unspecified Religiosity	25,487	0.05	0.22
Other Religions	25,487	0.01	0.09
Panel B: Peer group variables			
Peer group Size	996	25.64	21.64
High-ability share	996	0.35	0.24
Coethnic share	996	0.24	0.19
High-ability coethnic share	996	0.08	0.12
Low-ability coethnic share	996	0.15	0.15
Panel C: Course level			
All year grades (%)	1,061,905	67.83	9.61
Year One grades (%)	321,538	66.90	9.65
Year Two grades (%)	343,950	67.50	9.80
Year Three grades (%)	330,626	68.52	9.26

Notes: Data are from MUK and are restricted to students admitted to non-extension day majors at six colleges for 2009-2017, excluding 2010. A peer group comprises of students admitted to majors within a school major in the same year and assigned to the same dorm. Unspecified religion indicates whenever religious identities are not provided or entered as “Christian". Christianity is usually a correction of several or nondenominational religions in this context. Religion “Other" includes the smallest religions (where the count is less than 100 in the sample), such as Bahai, Jehovah's Witness, traditional religions, and Intambiro. Apart from age, Panel A variables are constructed to be binary.

2.5 Empirical Strategy

2.5.1 Mean Effects

Direct Effects of Coethnic and High-ability Peers

As aforementioned, a peer group refers to students admitted to majors in the same school f and assigned to dorm d in year t. For simplicity, I will index the peer group fdt with G in the estimation equations in this section. To estimate the direct effects of coethnic or high-ability peers on academic outcomes, I use a model that exploits variation in coethnic composition across peer groups within

Figure 2.4: Distribution of Coethnic Share across Peer Groups.

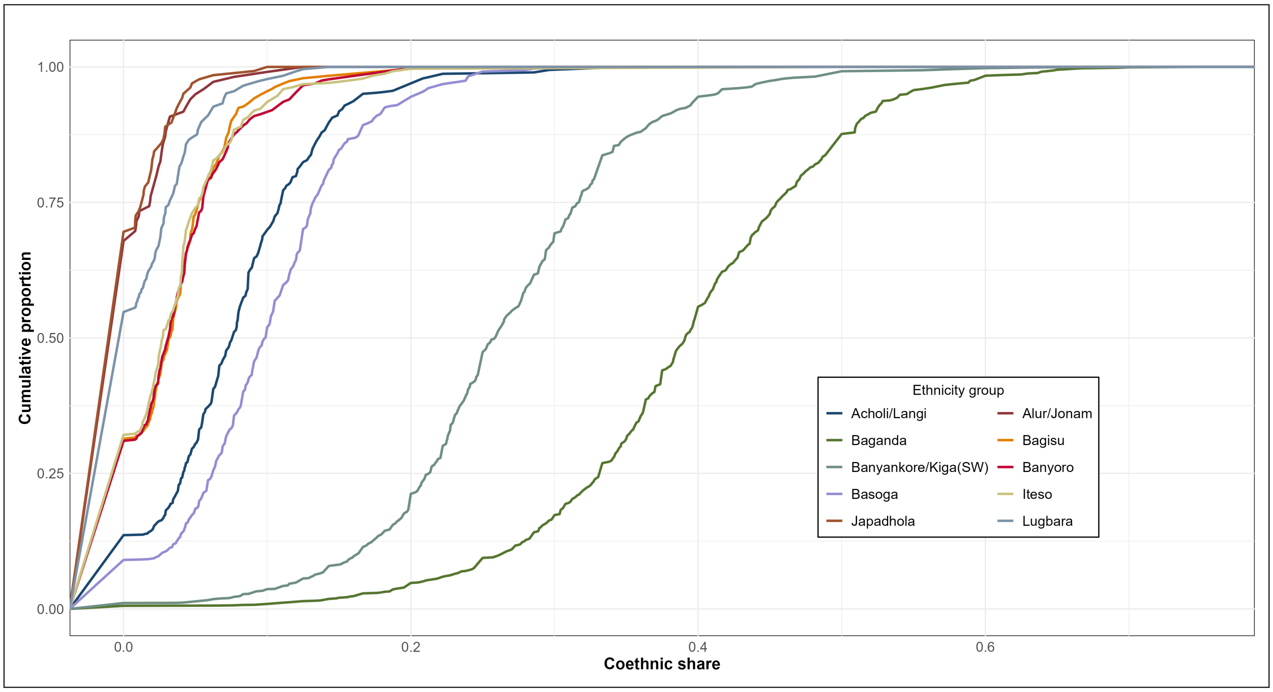

Notes: Data used to produce these distributions are from MUK and are restricted to students admitted to non-extension day majors from six colleges for 2009-2017, excluding 2010. Coethnic share is computed as the leave-me-out proportion of coethnic peers in a group. This figure plots the 10 largest ethnic groups by the total number of MUK students (out of the 16 total ethnic groups).

a year and year-to-year variation:

$$y_{ijcG} = \beta_0 + \phi_1 S^E_{iG} + \phi_2 S^H_{iG} + \beta_2 \mathbf{X}_{iG} + \beta_3 \bar{\mathbf{X}}_G + \delta_j + \alpha_c + \lambda_d + \theta_m + \gamma_s + \varepsilon_{ijcG}, \tag{2.4}$$

where y_{ijcG} is the first year percent grade that student i of ethnicity j and belonging to group G obtained in course c. S^E_{iG} is the probable coethnic share of in i's peer group defined in Section 2.4.2 in equation(2.3), and S^H_{iG} is the share of high-ability peers (coethnic and noncoethnics). The main estimation controls for δ_j, which is i's most probable ethnic group, $\mathbf{X}_{iG}$ is a vector of i's background characteristics and includes i's own ability, and $\bar{\mathbf{X}}_G = \frac{\sum_{\forall k \neq i} X_{iG}}{N_G - 1}$ is the vector of exogenous variables (the average background characteristics of i's peers, except high-ability). Additionally, α_c, λ_d, θ_m, and γ_s represent classroom, dorm, major, and high school subject combination fixed effects (FE). Lastly, ε_{ijcG} is the error term. I cluster standard errors at the peer group to account for the potential error correlation across individuals in a group.

The coefficients of interest are ϕ_1, which captures the effect of attending lectures and potentially living with coethnic peers in this setting, and ϕ_2, which captures the effect of attending class and potentially living with high-ability peers irrespective of their ethnicity. I take several steps to ensure that ϕ_1 and ϕ_2 are unbiased. I control for several FE to deal with bias arising from correlated shocks.

First, correlated shock in this setting may arise from differences across classrooms. Therefore, I include classroom FE to control for unobserved differences in courses, such as performance, instructor effects, and classroom diversity. In addition, classroom FE should control for differences in major by year since the student's major and course list are determined at the time of admission. Nevertheless, students may take courses with peers admitted to majors outside their schools if cohort sizes are small and major course requirements are related. This implies nonrandom exposure to coethnic peers because of the systematic differences in the share of some ethnicities in the MUK sample and general population. Thus, α_{ct} also controls for this systematic difference in ethnic exposure across students in addition to controlling ethnicity FE. Relatedly, I include major FE, θ_m to control for differences between majors. Each regression will control for ethnicity, major, and classroom FE at the minimum in the results section.

Second, I control for dorm FE to control for factors, such as renovation and dorm conditions, that might affect academic performance. In addition, cultures are different across dormitories. For example, Ricart-Huguet and Paluck (2023) show that cultures, such as outgoing and academic mindedness, are different across MUK dorms to the extent that they affect interpersonal outcomes.

Third, although evaluated at the same cutoffs, students entering the same major may occasionally take different subject combinations during upper high school. Therefore, I include high school subject group FE, γ_s, to capture the differences in types of incoming students. When computing high school weighted GPA, each major has different requirements to capture incoming students' academic preparedness. Take Bachelor of Commerce, for example, the required HS subjects are math and economics, but students who take one of the two and those who take both can qualify if they perform above the cutoffs. Students graduating with math and economics have a higher perceived potential for success in Bachelor of Commerce classes than those graduating with one of the two subjects. Therefore, controlling γ_s captures the unobserved differences in academic preparedness across students in the same major.

Concerning self-selection, dorm assignment is random, and each student's course list and classmates are predetermined before entry at the time of admission, as mentioned in Section 2.3.1. Two lines of concern can be made for potential sources of self-selection. First, although dorm assignment is random, the on-campus residence may be biased to STEM students admitted through the merit scholarship. This is only statistically meaningful if merit scholarships are correlated with ethnicity and if dorm assignment was not random.

Correlation between ethnicity and obtaining a scholarship in a STEM major is possible if the top secondary schools are concentrated in one ethnic region, where students from that region graduate with the highest A-level scores to qualify for the merit scholarship. As Panel B of Figure 2.2 shows, the student population is biased towards the two largest ethnic groups. Coincidentally, the most elite secondary schools in the country belong to these regions because of historical reasons. Nevertheless, this is not an issue, as dorms and majors do not have ethnic quotas and equation (A1) controls for i's probable ethnic group, which controls for differences in the levels of stratification. Also, when I regress merit ethnicity fixed effects, I find the explained variation is less than 1%.

Additionally, students may select into majors by manipulating the rank of their choices. This might cause selection into majors even though an organization separate from the university handles admissions and even though obtaining admission is quasi-experimental. This is possible since the ranking of program cutoffs does not change from one year to another, although actual cutoffs may change. This is not a concern as a peer group of classmates who potentially live together, and dorm assignment is random.

As a test, I provide balance tests in Table 2.2, which presents evidence against selection. Each column is an independent estimation similar to specification (A1). I run these regressions at the aggregated to the student level (not course level). Each pre-university characteristic is regressed against the coethnic share in Panel A, while in Panel B, each pre-university characteristic is regressed against the high-ability coethnic share. Panel B also controls for a student's ability. Additionally, I regress the share of coethnic or high-ability peers on all the pre-university characteristics and report the estimates and the F stat in Appendix Table A4. The correlation between pre-university characteristics and the primary variable of interest would be high and significant if nonrandom sorting into peer groups existed.

From Table 2.2, the correlation between each student's characteristics and the share of

Table 2.2: Evidence against Selection

	(1) Age	(2) Anglican	(3) Catholic	(4) Muslim	(5) Pentecostal	(6) SDA	(7) High Ethnic Salience	(8) Other Religion	(9) High-ability
Panel A: Coethnic share as the independent variable									
Coethnic share	-0.00	-0.01	0.02	0.01	-0.02	-0.01	0.04	0.00	0.02
	(0.11)	(0.04)	(0.03)	(0.02)	(0.02)	(0.01)	(0.04)	(0.01)	(0.03)
R-squared	0.13	0.07	0.05	0.27	0.03	0.03	0.11	0.06	0.17
N	25,487	25,487	25,487	25,487	25,487	25,487	25,487	25,487	25,487
Panel B: High-ability share as the independent variable									
High-ability share	0.06	-0.00	-0.02	-0.01	0.02	-0.01	0.03	0.01	
	(0.08)	(0.03)	(0.03)	(0.01)	(0.01)	(0.01)	(0.03)	(0.01)	
R-squared	0.13	0.07	0.05	0.27	0.03	0.03	0.13	0.06	
N	25,487	25,487	25,487	25,487	25,487	25,487	25,487	25,487	25,487

Notes: Data are from MUK and are restricted to students admitted to non-extension day majors from six colleges for 2009-2017, excluding 2010. Each column is an independent regression that regresses a pre-university characteristic against the share of coethnics. All regressions include school-by-year (not classroom), ethnicity, and hall FEs. Also, all regressions in panel B control for own ability. Standard errors are in parentheses and clustered at the peer group level.
*p<0.1, **p<0.5, ***p<0.01

coethnic peers (Panel A) or the share of high-ability peers (Panel B) is practically zero and not significant in all regressions, providing evidence that students are not selecting into peer groups. In addition, the F stats from Appendix Table A4 are small, 0.84 and 2.15 when the share of coethnic (column one) or share of high-ability (column two) peers are regressed against all the pre-university characteristics, respectively. This indicates that the results presented in this paper are unlikely to be biased because of nonrandom sorting.

Interpreting the Magnitude of ϕ_1 and ϕ_2

As Section 2.3.2 describes, not all students assigned to a dorm end up residing in their assigned dorm due to capacity constraints. Living off campus does not, however, exclude a student from dorm-based peer groups; it just alters the nature and frequency of interactions. Given my peer group definition, consider two types of students based on the extent of interactions with others in a given peer group: fully compliant and partially compliant. Fully compliant students live in their assigned dorms and can thus interact as dorm residents with other fully compliant peers and, in other ways, with their partially compliant peers. Partially compliant students live off-campus and, therefore, do not interact as dorm residents with others in the peer group I construct for them. Both types of students are likely to interact daily (within and across each type) in classes and study groups.

If I observed residents, I could estimate the local average treatment effect of coethnic and high-ability peers using dorm assignment as an instrument for dorm residence to account for endogenous dorm residency. Since I do not observe residence, ϕ_1 and ϕ_2 in equation (A1) are effectively reduced-form peer effects estimates based on dorm assignment. These reduced-form effects are a data-weighted average of the peer effects for fully compliant peers and partially compliant peers.[6]

I expect the reduced-form peer effects in Equation (A1) to be less than the local average treatment effect of coethnic and high-ability peers. Also, the existence of partially compliant peers will likely attenuate peer effects. To illustrate the logic behind this claim, consider a parallel with Carrell, Fullerton, and West (2009), who use dorm floors to reconstruct peer groups at the Air Force. Their empirical setting allows them to construct pseudo-peer groups that span the relevant peer group (squadrons). Interactions are expected to still exist within these pseudo-peer groups but at a reduced rate than the true squadron-based peer groups. Although these pseudo-peer groups comprised 66.6% of peers from squadrons, the presence of peers with whom students interact less frequently attenuates estimated peer effects. Analogously, in the MUK context, I expect that reduced-form peer effects based on dorm assignment and, therefore, including both resident and non-resident students in peer groups will underestimate the true peer effects operating in this setting.

Coethnic and High-ability Interaction Effects

Although evidence in the literature is mixed, the effect of the high-ability peers (ϕ_2) is expected to be positive. Classrooms in the setting are relatively large, and services, such as office hours, do not exist, making peer-to-peer learning important. The sign of the coethnic share coefficient (ϕ_1) is largely ambiguous. The coefficient ϕ_1 could be zero on average, although it could be positive or negative for several reasons.

Bayer et al. (2020) show that minority students in introductory economics classes report a lower sense of belonging than non-minority students, and some studies (e.g., Walton and Cohen

[6] Even knowing the overall proportion of each cohort residing in dorms (i.e., residence compliance) does not necessarily solve this issue as I cannot use this compliance rate as the first stage to inversely weight reduced-form in the equation proposed in Bloom (1984) and as equation (A4b) in Appendix Section 2.8.3 without additional (strong) assumptions.

2011) show that interventions to increase a sense of belonging improve academic outcomes for students. Additionally, some students might suffer from imposter syndrome exacerbating their sense of belonging. Thus, having coethnic peers might be significant for some student as it may increase a sense of belonging during student interactions.

On the other hand, students of shared ethnicity may gravitate toward one another for cultural reasons, such as language, traditions, and beliefs. These homophilous tendencies and coethnic bias during interactions in diverse societies might lead to ethnic-based sorting into study and friendship groups. In this case, the effect of coethnic peers on academic performance will be indirectly through high-ability coethnic peers. For example, it might be detrimental for coethnic peers to isolate if they are, on average, low-ability compared to noncoethnic peers. A low-ability student might benefit from a higher share of high-ability than a higher coethnic share in a peer group.

To capture the effect of the pre-university academic quality of coethnics, I use an equation similar to equation (A1).

$$y_{ijcG} = \beta_0 + \phi_1 S_{iG}^{EL} + \phi_2 S_{iG}^{EH} + \phi_3 S_{iG}^{E'H} + \beta_2 \mathbf{X}_{iG} + \beta_3 \bar{\mathbf{X}}_G + \delta_j + \alpha_{ct} + \lambda_d + \theta_m + \gamma_s + \varepsilon_{ijcG}, \tag{2.5}$$

where S_{iG}^{EH}, S_{iG}^{EL} and $S_{iG}^{E'H}$ are the probable shares of high-ability coethnic, low-ability coethnic, and high-ability noncoethnic peers, respectively. All other terms are the same as those in equation (A1).

Therefore, the setting provides four sources of variation of interest in the share of peers that is: (A) high-ability and coethnic; (B) low-ability and coethnic; (C) high-ability and noncoethnic; and (D) low-ability and noncoethnic. Coefficient ϕ_1 in equation (A1) captures the average effect of (A) and (B), while ϕ_2 captures the average effect of (A) and (C). If students are sorting on ethnicity when forming study groups, (A) should matter than (C). In such cases, we can think of the coethnic peers operating indirectly through high-ability ethnic peers.

2.5.2 Heterogeneous Peer Effects

To estimate heterogeneous effects, I estimate equation (A1), including the interaction of the two treatments with the dummy variable that captures the heterogeneous dimension listed below.

(2.6)

$$y_{ijcG} = \beta_0 + \phi_1 S^E_{iG} + \phi_2 S^H_{iG} + \varphi_1 S^E_{iG} \times d_i + \varphi_2 S^H_{iG} \times d_i + \beta_2 \mathbf{X}_{iG} + \beta_3 \bar{\mathbf{X}}_G + \delta_j + \alpha_c + \lambda_d + \theta_m + \gamma_s + \varepsilon_{ijcG},$$

where d_i can be gender, ability, or assumed level of ethnicity salience, while φ_1 and φ_2 are the differential impacts on d_i of coethnic and high-ability share, respectively. All other terms are the same as those in equation (A1).

If coethnic and high-ability peers matter for academic performance, the average effects may be conceptually different depending on dimensions, such as the size of each ethnicity at the university and level of prior exposure to noncoethnic Ugandans. If increasing a sense of belonging is a channel through which coethnic peers might work, then the effect of coethnic peers could be zero for large ethnic groups who are less likely to suffer a low sense of belonging. However, coethnic peers may matter positively for small ethnic groups with limited exposure to different ethnicities prior to University.

Also, coethnic peer effects in the presence of high ethnic heterogeneity may also matter due to inter-ethnic impacts. For instance, there are ethnic groups that share values with other groups or portray less in-group bias. In such cases, fewer co-ethnic peers may not matter as those students would easily integrate with other ethnicities. More broadly, if inter-ethnic uncongenial relationships exist in Ugandan societies, they could spill over into schools, creating 'bad' peers. Nevertheless, this is unlikely in Uganda, as inter-ethnic tensions are not that common. The analysis will, therefore, explore heterogeneity in other dimensions.

Differential Effects by Gender

Several studies report differential peer effects on academic and non-academic outcomes by gender in several settings. For example, Carrell and Hoekstra (2010) study peer effects of kids exposed to domestic violence on test scores and disciplinary incidents in a classroom and find that peer effects are significant and stronger for boys, not girls. Additionally, Stinebrickner and Stinebrickner (2006) use HS GPA to study peer effects on sudy habbits and academic perforamcne at Berea College and find that HS GPA captures the effect on study habits and significant peer effects on girls. More recently, using a field experiment at a public school in Peru, Zárate (2023) finds low peer effects on academic outcomes but stronger on social skills, such as network connectivity, and psychological measures of social skills, such as altruism, which vary by gender.

Given the coethnic bias reported in the literature, it is likely that friendships are formed along ethnic lines. For example, using a setting in SSA similar to Uganda, Salmon-Letelier (2022) report that friendship networks form along ethnicity or religion lines in Nigeria's state and federal unity schools, respectively. If such homophily exists, it may create differential impacts by gender since friendship groups overlap with study groups.

There is also long-standing anthropological literature, such as de la Cadena (1995) exploring ethnicity and gender that finds women are more ethnic than men in the community of Cusco. Studying how information affects homophily, Gallen and Wasserman (2023), finds that women portray homophile tendencies more than men in an online college mentoring platform. Also, Jackson et al. (2022) track university students' friendships and study partnerships in their Caltech Cohort study and find assortative homophily by gender and ethnicity exists and persists substantially over time among friendship and study groups.

Differential Effects by Ethnic Salience

Having a coethnic in a peer group might be useful for students with high ethnic salience due to migrating from their home regions to attend university and experiencing a "diversity shock" when they arrive at the campus. Migrating from one's ethnic region to attend a university located in a different ethnic region with different cultures could cause immigrant students to be aware of their own ethnic identity, leading to greater attachment to their own ethnicities. This is the phenomenon in Okunogbe (2018), who finds greater ethnic pride among Nigerian youth randomly assigned to serve in a region where the ethnic majority is different from their own ethnicity. These hypotheses also align with the psychology literature on social identity, which suggests that the salience of one's ethnic identity increases when one migrates away from one's native region.

In addition, such students could face social isolation as they encounter cultural barriers, which may increase their stress levels and contribute to a lack of sense of belonging. Moreover, students from certain ethnicities may experience discrimination from other groups, leading them to isolate themselves.[7] These students are forced to navigate a new learning environment where

[7] Another potential reason for the isolation of certain ethnicities is inter-ethnic conflicts and competition spilling

classrooms are more diverse than their high schools. Yet, several studies report generally low trust levels in addition to high in-group bias in highly ethnically diverse societies. Having a high proportion of coethnic peers in their peer group can be beneficial for students with high ethnic salience, especially if they belong to small ethnic groups.

Differential Effects by Degree Type

Studies on post-secondary education have reported differences by subject type. For example, Carrell, Fullerton, and West (2009) find peer effects are stronger in math and science courses, smaller in social sciences, and absent in foreign languages and physical education at the Air Force Academy. Studying peer effects from the field of study at an Italian university, Brunello, De Paola, and Scoppa (2010) find that peer effects are stronger in the 'hard' sciences (engineering, math, and natural sciences) but absent in social sciences and humanities. I do not observe course names, but I observe the course code (e.g., STA101) and the degree type. MUK offers degrees in either arts or sciences. Arts degree comprises a wide range of degrees, such as business-related, social sciences, and humanities, and so do science degrees.

There are other reasons in this context to anticipate why peer effects might differ by degree type. For example, classrooms in arts degrees may differ from those in science classes, as they are, on average, larger. Additionally, the proportion of high-ability peers in arts degrees is smaller due to the design of the national merit scholarship scheme. Generally, larger class sizes would reduce interaction with the professor by increasing the student-to-teacher ratio. Since student-teacher interactions outside the classroom are limited in this setting, class size effects are more likely to manifest through peer effects. Also, large classrooms might increase the need for a sense of belonging and may reduce intimate cross-cultural interactions among students. It is easier to sort based on ethnicity as the probability of the coethnic presence of ethnicity is high.

Differential Effects by Ability

Heterogeneous peer effects by a student's own ability and peers' average ability have been shown to exist in the literature. For example, Zimmerman (2003) finds students in the middle of the verbal SAT distribution have negative peer effects from low-ability roommates. Also, Carrell, Fullerton,

into classrooms, although ethnic conflicts are not common in this setting.

Table 2.3: Mean Effects in Year One: Coethnic vs High-ability Share

	(1)	(2)	(3)	(4)
Coethnic share	1.054**	1.064**	1.046**	0.936**
	(0.47)	(0.47)	(0.47)	(0.47)
High-ability share	0.799***	0.833***	0.848***	0.735**
	(0.29)	(0.29)	(0.29)	(0.28)
High-ability	3.790***	3.792***	3.791***	3.787***
	(0.09)	(0.09)	(0.09)	(0.09)
R-squared	0.326	0.326	0.328	0.328
N	321,452	321,452	321,452	321,452
Dorm FE	No	Yes	Yes	Yes
Individual Controls	No	No	Yes	Yes
Group Controls	No	No	No	Yes

Notes: Data are from MUK and are restricted to students admitted to non-extension day majors from six colleges for 2009-2017, excluding 2010. A peer group comprises students admitted to majors within the same school assigned to the same dorm. Each column is an independent regression, but the outcome is the course grades in all the specifications. The differences between each specification are indicated at the bottom and come from the controls. All regressions control for own ethnicity, gender, own ability, major FE, and HS subject combination FE, but gender drops out (2)-(4) since dorms are single-sex. Individual controls include age, religious indicators, and graduating from the district of origin. Group controls include the leave-me-out averages of individual controls in addition to peer group size. SEs are parentheses and are clustered at the peer group level.
*p<0.1, **p<0.5, ***p<0.01

and West (2009) find suggestive evidence of non-linearity peer effects. Verbal SAT peer effects are strong for students in the bottom third of the distribution. Given reduced student-teacher interaction in this setting, peer effects may exist through study partnership channels, especially for low-ability students.

2.6 Results

2.6.1 Mean Effects

Table 2.3 estimates various specifications of equation (A1). All specifications control for ethnicity as described in section 2.5, ability, gender, student's major, and classroom and HS subject combination fixed effects. The difference between specifications is shown at the bottom of each column. It comes from controls in each regression, as I begin with a simple regression and progressively add more controls. Since I do not know the residence status of the students, the coefficients reported in this should be interpreted as intent to treat effects.

Given no evidence of selection, as reported in Section 2.5, we do not expect the coefficients to change significantly as we move from column (1) to (4). The table shows that the share of coethnic and high-ability matters significantly for academic performance. The effect of coethnic share is stable at around one percentage point (pp) , while that of high-ability peers is around 0.8pp. Adding dorm fixed effects and individual and group controls does not alter the effects.

The results in specification From specification (4) imply that adding five coethnic peers to a typical peer group of size 25 (corresponding to the average group size) increases a student's performance by 0.19pp ($5/25 \times 0.936$). This effect is equivalent to 0.02 standard deviations in a student's performance in the first year. The effect of high-ability share is 0.735, which implies that adding two more high-ability peers to a typical group of size 25 increases a student's performance by 0.15pp ($5/25 \times .735$), which is also about 0.02 standard deviations change in a student's performance.

For context, adding five coethnic peers to a typical peer group of size 25 (corresponding average size as Table 2.1 shows) is equivalent to moving from the group where the number of coethnic peers corresponds to the twenty-fifth percentile to a group where the number of coethnic peers corresponds to the seventy-fifth percentile.[8] For simplicity, I will interpret the results as the effect of adding either five coethnic or high-ability peers to a group of 25.

Table 2.3 also shows the effect of own ability is much larger than the effect of coethnic and high-ability share. The table shows that high-ability students perform about four percentage points higher than low-ability peers. This difference is large as it corresponds to a change in grade that would move a student whose first-year grade is equal to the average from the second-class lower (Fairly Good) performance range to a second-upper (Good) range. The average of grades reported in Table 2.1 is equivalent to second class lower degree type in this setting.

Results show positive and direct peer effects from a higher share of high-ability (irrespective of ethnicity) and coethnic peers. However, it is likely that the high ability of coethnic peers might matter, while high-ability noncoethnic peers do not. To test this, I break down the treatment variables into four: (A) high-ability coethnic peers, (B) low-ability coethnic peers, (C) high-ability noncoethnic peers, and (D) low-ability noncoethnic peers and compute the share of each as described

[8]It is important to note that distribution of high-ability peers may be different from that of coethnic peers in a group. I use a marginal change of 5 peers in a typical group size for simplicity of interpretation.

Table 2.4: The Effect of High-ability Coethnic and High-ability Noncoethnic peers on Academic Performance in Year One.

	(1)	(2)	(3)	(4)
(A) High-ability coethnic share	0.928*	0.980*	1.053**	0.865*
	(0.52)	(0.52)	(0.51)	(0.50)
(B) Low-ability coethnic share	0.686*	0.692*	0.696*	0.648
	(0.41)	(0.41)	(0.41)	(0.41)
(C) High-ability noncoethnic share	0.963***	0.992***	0.991***	0.875***
	(0.32)	(0.32)	(0.32)	(0.32)
R-squared	0.326	0.326	0.328	0.328
N	321,375	321,375	321,375	321,375
Dorm FE	No	Yes	Yes	Yes
Individual controls	No	No	Yes	Yes
Group controls	No	No	No	Yes

Notes: Data are from MUK and are restricted to students admitted to non-extension day majors from six colleges for 2009-2017, excluding 2010. A peer group comprises students admitted to majors within the same school, and assigned to the same dorm. Each column is an independent regression, but the outcome is the course grades in all the specifications. The differences between each specification are indicated at the bottom and come from the controls. All regressions control for own ethnicity, gender, own ability, major FE, and HS subject combination FE, but gender drops out (2)-(4) since dorms are single-sex. Individual controls include age, religious indicators, and graduating from the district of origin. Group controls include the leave-me-out averages of individual controls in addition to peer group size. SEs are parentheses and are clustered at the peer group level.
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

in Section 2.5. If high-ability coethnic peers matter while high-ability noncoethnic peers do not, (A) should be positive and significant while (C) should not, or at least (A) should be larger than (C), indicating that coethnic peers matter indirectly through high-ability peers.

These results in Table 2.4 largely follow the pattern observed in Table 2.3. When running these regressions, I exclude category (D). Therefore, the reported coefficients should be interpreted relative to that reference group. From the preferred specification (4), adding five high-ability coethnic or noncoethnic peers to a group of size 25 increases a student's course grade by 0.18pp relative to low-ability noncoethnic peers, which is equivalent to 0.02 standard deviations. The same table also shows that low-ability coethnic peers have a large and positive effect on academic performance. However, it is imprecise.

Lastly, estimates in both Table 2.4 and Table 2.3 don't change or change very little when I add dorm FE, individual controls, and group controls. This is consistent with exogenous assignment into peer groups and that correlated shocks are less likely to drive results reported in this paper. The results do not suggest that high-ability coethnic peers matter more than high-ability noncoethnic peers in the first year, which is contrary to my prior. Taken together, these results show that, on average, high-ability peers directly and significantly impact every student's grades regardless of their ethnicity. Also, these results show that coethnic peers have a positive effect on grades, although they suggest high-ability coethnic matter more than low-ability coethnic peers.

2.6.2 Persistence of Mean Effects

All the results presented thus far focus on student performance during their first year at MUK. By extending the analysis to the subsequent years of their education, I test for the persistence of these peer effects. If peer effects from social networks persist as a student advances throughout their college career, then the effects of high-ability and coethnic peers observed in Section 2.6.1 should also be evident in the follow-on years. Since selection into courses is limited, this setting allows me to explore the persistence of peer effects. I estimate specification (A2) separately for each of the three academic years of undergraduate education at MUK and present the results in Table 2.5. Panel A compares the effect of coethnic share to that of high-ability share, while Panel B compares high-ability coethnic to noncoethnic peers.

Column (1) restates the results in Section 2.6.1 to facilitate comparison. The effect of coethnic share persists into the second year, but it is almost half of the magnitude of the first year's effect in the third year. That is, the effect of adding five coethnic peers into a peer group of size 25 is 0.10pp in the third year, which is not statistically different from zero and is almost half of the effect observed in the first year, as reported in Table 2.3. In comparison, the effect of high-ability share persists and even increases in the third year. From Panel A, adding two high-ability peers to a group of 25 increases a student's performance by 0.22pp in the third year Yet, the same change increases student performance by 0.15pp in the first year. Thus, the effect of high-ability peers in the third year is 1.5 times the effect observed in the first year, and that of coethnic peers is one-half of what is observed in the first year.

Panel B shows results similar to those in Panel A. Relative to low-ability noncoethnic peers, the effect of high-ability coethnic and noncoethnic in the third year is positive and significant. In contrast, that of low-ability coethnic peers in the third year is not significant and is about 25% lower than the effect observed in the first year. Additionally, the effects of high-ability peers (coethnic and noncoethnic) relative to low-ability peers increase from the first to the third year.

The results in Table 2.5 show that the role of shared identity, if not coupled with ability, falls as time goes on. However, the effect of high-ability peers rises as time goes on, although the effect of high-ability coethnic peers increases more than that of high-ability noncoethnic peers. The results are suggestive of evolving study groups or social networks.[9] For example, students might form stronger study bonds with high-ability peers as time goes on.

2.6.3 Heterogeneous Peer Effects

Results in Section 2.6.1 show that, on average, going to school and potentially living with coethnic and high ability affects academic performance, but the effect of the coethnic share falls as time goes on. I now turn to see whether there are differential impacts on different dimensions mentioned in Section 2.5.2, as such differential effects might shed light on some mechanisms.

Differential Impacts by Gender

Table 2.6 presents the differential effects by gender. As dorms are single-sex, I control for gender instead of dorm FE. Controlling for dorm FE does not change the results, but gender drops out. The table also represents results across the years and p-values corresponding to testing the significance of the treatment variables for girls: coethnic share ($\phi_1 + \varphi_1$) and high-ability share ($\phi_2 + \varphi_2$) in equation (2.6). These results reveal several patterns.

First, girls perform lower than boys by 0.743pp, significant in the first year, but they perform better than boys by 1.061pp in the third year. The difference in performance between boys and girls is not significant in the second year.

Second, the coethnic share is positive but not significant for both boys in the first and third years, but it is marginally significant in the second year. Although economically meaningful, the differential impact by gender is not significant across years. These results suggest that, unlike boys, girls might be benefiting from a higher coethnic share. The p-value testing the sum of the

[9] I will explore this mechanism when I run surveys later. It is possible that students hang out with coethnic peers at the start of their university career because it is more organic. However, as time goes on, networks may evolve as they learn which of their peers are high-ability (coethnic or noncoethnic). They might form stronger networks with high-ability coethnic peers, weaker networks with low-ability coethnic peers, and somewhat strong networks with high-ability noncoethnic peers, as high-ability peers may be perceived as more beneficial for academic performance. The survey will ask students about their study and friendship groups throughout their undergrad career to see if they are constant or changing overtime.

Table 2.5: Persistence of Mean Effects in Follow-up Years

	Year One	Year Two	Year Three
Panel A: Coethnic vs High-ability			
Coethnic share	0.936**	1.136**	0.491
	(0.47)	(0.48)	(0.45)
High-ability share	0.735**	0.839***	1.101***
	(0.28)	(0.28)	(0.28)
R-squared	0.328	0.384	0.380
N	321,375	343,761	330,158
Panel B: High-ability (Coethnic vs Noncoethnic)			
High-ability coethnic share	0.867*	1.201**	1.347***
	(0.50)	(0.52)	(0.50)
Low-ability coethnic share	0.639	0.775*	0.476
	(0.41)	(0.41)	(0.39)
High-ability noncoethnic share	0.879***	0.951***	1.166***
	(0.32)	(0.32)	(0.32)
R-squared	0.326	0.383	0.378
N	310,867	333,187	320,752
Dorm FE	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes
Group controls	Yes	Yes	Yes

Notes: Data are from MUK and are restricted to students admitted to non-extension day majors from six colleges for 2009-2017, excluding 2010. A peer group comprises students admitted to majors within the same school and assigned to the same dorm. Each column is an independent regression, but the outcome is course grades in all regressions. All regressions control for own ethnicity, own ability and major, HS subject combination, and classroom FE. Individual controls include age, religious indicators, and graduating from the district of origin. Group controls include the leave-me-out averages of individual controls in addition to peer group size. SEs are parentheses and are clustered at the peer group level.
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

significance of coethnic share and the interaction of the coethnic share and female dummy is largely significant in the first and second years and marginally significant in the third year. These results imply that adding five coethnic peers in a group of 25 increases academic performance for boys by 0.13pp, which is equivalent to a 0.01 standard deviation change in academic performance in the first year. On the contrary, the same in coethnic peers would increase girl's performance by 0.25PP, equivalent to 0.03 standard deviation in the first year. Thus, the effect of coethnic share on boys is about 30% lower than the average effect in the first year, yet the effect on girls is about 30% larger than the average effect observed in Table 2.3.

Third, the differential effect of high ability by gender is small, insignificant, and sometimes

Table 2.6: Differential Effect by Gender: Coethnic vs High-ability Share

	Year One	Year Two	Year Three
Coethnic share	0.630	0.928*	0.136
	(0.55)	(0.56)	(0.53)
High-ability share	0.661**	0.821***	1.101***
	(0.31)	(0.31)	(0.31)
Female	-0.743***	0.243	1.061***
	(0.21)	(0.21)	(0.21)
Coethnic share × Female	0.605	0.398	0.668
	(0.48)	(0.49)	(0.46)
High-ability share × Female	0.136	-0.021	-0.074
	(0.42)	(0.42)	(0.41)
p-val Coethnic share: Female	0.015	0.001	0.097
p-val High-ability share: Female	0.059	0.053	0.013
R-squared	0.328	0.384	0.380
N	321,375	343,761	330,158
Dorm FE	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes
Group controls	Yes	Yes	Yes

Notes: Data are from MUK and are restricted to students admitted to non-extension day majors from six colleges for 2009-2017, excluding 2010. A peer group comprises students admitted to majors within the same school and assigned to the same dorm. Each column is an independent regression, but the outcome is course grades in all regressions. All regressions control for own ethnicity, major, HS subject combination, and classroom FE. Individual controls include age, religious indicators, and graduating from the district of origin. Group controls include the leave-me-out averages of individual controls in addition to peer group size. SEs are parentheses and are clustered at the peer group level. The table also reports p-values for coethnic share and high-ability share of female students. These tests correspond to $\phi_1 + \varphi_1$ and $\phi_2 + \varphi_2$ in equation (2.6).
* p<0.10, ** p<0.05, *** p<0.01

negative. Still, the share of high-ability peers has a positive and significant effect on boys and girls. As the average effect reported in Table 2.3, the share of high-ability persists and increases into the third year for both boys and girls.

Differential Impacts by Ability

Table 2.7 shows the differential effects by own ability across the years. The table also reports results across the years and p-values corresponding to testing the significance of the treatment variables for high-ability: coethnic share ($\phi_1+\varphi_1$) and high-ability share ($\phi_2+\varphi_2$) in equation (2.6). High-ability students perform higher than low-ability peers, as Section 2.6.1 already reported. The results reveal several other patterns.

First, the effect of coethnic share on the academic performance of low-ability students is not significant across all the years and is negative in the third year. On the other hand, the effect of high-ability share on low-ability students is positive and significant across all the years and even higher in the third year than in the first year. For example, adding five high-ability peers to a group of 25 increases a slow student's performance by 0.14pp and 0.25pp in the first and third year, respectively. Although the differential impact by ability is not significant across all the years, it is negative and economically meaningful in the third year, implying that high-ability peers have a larger effect on low-ability students than high-ability students.

Second, the differential impact of coethnic share is large and significant in Year One and largely persists into Year Three. Table 2.7 shows adding five coethnic peers to a group of 25 increases the performance of high-ability students by 0.51pp, which is about 3.5 times the average effect reported in Table 2.3. This effect is about 0.05 standard deviation of the first year's performance. This differential impact of coethnic share on high-ability students persists significantly into the third year, albeit at a reduced magnitude.

Differential Impacts by Degree Type

I estimate the differential effect by degree type and report it in Table 2.8. Since I control for the Major FE, the arts major dummy drops out of the regressions. Although not shown, the results change significantly when I control for the degree type dummy instead of major FE changes. The table also reports results across the years and p-values corresponding to testing the significance of the treatment variables for arts majors: coethnic share ($\phi_1 + \varphi_1$) and high-ability share ($\phi_2 + \varphi_2$) in equation (2.6).

This table shows the differential impact of high-ability share by degree type is large and significant across all the years and more than doubles from the first to third year. The results show that adding five high-ability peers to a group of 25 increases a student in the art's major performance by 0.33pp in the first year, which is almost 1.5 times the average effect reported in Table 2.3. Moreover, this effect increases to 0.362pp in the third year, which is 3.3 times the effect reported 2.5. The effect of high-ability share on a student who is a degree major in the third year

Table 2.7: Differential Effect by Ability Type: Coethnic vs High-ability Share

	Year One	Year Two	Year Three
Coethnic share	0.064	0.274	-0.126
	(0.48)	(0.48)	(0.46)
High-ability share	0.716**	0.911***	1.234***
	(0.31)	(0.32)	(0.30)
High-ability	3.259***	3.013***	2.816***
	(0.20)	(0.20)	(0.19)
Coethnic share × High-ability	2.467***	2.513***	1.846***
	(0.51)	(0.51)	(0.48)
High-ability share × High-ability	-0.015	-0.269	-0.429
	(0.42)	(0.44)	(0.41)
p-val Coethnic share: High-ability	0.000	0.000	0.003
p-val High-ability share: High-ability	0.082	0.115	0.046
R-squared	0.328	0.384	0.380
N	321,452	343,840	330,236
Dorm FE	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes
Group controls	Yes	Yes	Yes

Notes: Data are from MUK and are restricted to students admitted to non-extension day majors from six colleges for 2009-2017, excluding 2010. A peer group comprises students admitted to majors within the same school and assigned to the same dorm. Each column is an independent regression, but the outcome is course grades in all regressions. All regressions control for own ethnicity, own ability, and major, HS subject combination, and classroom FE in addition to individual and group controls. Individual controls include age, religious indicators, and graduating from the district of origin. Group controls include the leave-me-out averages of individual controls in addition to peer group size. SEs are parentheses and are clustered at the peer group level. The table also reports p-values for coethnic share and high-ability share of high-ability. These tests correspond to $\phi_1 + \varphi_1$ and $\phi_2 + \varphi_2$ in equation (2.6).
* p<0.10, ** p<0.05, *** p<0.01

is very large, as it corresponds to 0.08 standard of academic year in the third year.

Lastly, the table also shows the differential impact of coethnic share by degree type is significant in the first year but not in the second and third years. Adding five coethnic peers to a group of 25 increases the academic performance of a student in the arts degree by 0.22pp more than for a student in the science majors.

Table 2.8: Differential Effect by Degree Type: Coethnic vs High-ability Share

	Year One	Year Two	Year Three
Coethnic share	0.586	1.101**	0.650
	(0.51)	(0.52)	(0.49)
High-ability share	0.455	0.384	0.034
	(0.36)	(0.37)	(0.36)
Coethnic share × Arts degree	1.080**	0.188	-0.074
	(0.49)	(0.49)	(0.46)
High-ability share × Arts degree	0.975*	1.321**	2.979***
	(0.59)	(0.60)	(0.57)
p-val Coethnic share: Arts degree	0.006	0.017	0.283
p-val High-ability share: Arts degree	0.000	0.000	0.000
R-squared	0.328	0.384	0.380
N	321,375	343,761	330,158
Dorm FE	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes
Group controls	Yes	Yes	Yes

Notes: Data are from MUK and are restricted to students admitted to non-extension day majors from six colleges for 2009-2017, excluding 2010. A peer group comprises students admitted to majors within the same school and assigned to the same dorm. Each column is an independent regression, but the outcome is course grades in all regressions. All regressions control for own ethnicity, ability, major and HS subject combination, and classroom FE. Individual controls include age, religious indicators, and graduating from the district of origin. Group controls include the leave-me-out averages of individual controls in addition to peer group size. SEs are parentheses and are clustered at the peer group level. The table also reports p-values for coethnic share and high-ability share of arts degree. These tests correspond to $\phi_1 + \varphi_1$ and $\phi_2 + \varphi_2$ in equation (2.6).
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Differential Impacts by Ethnic Salience

I proxy high 'ethnic salience' using a dummy variable equal to one if a student graduated high school from a district of origin and zero otherwise.[10] Since most Ugandan districts are ethnically segregated, these students have generally had much less exposure to other ethnicities prior to enrolling at MUK than their peers of the same ethnicity who graduated high school outside their district of origin (e.g., as a boarding student in or near Kampala). I estimate the differential impacts by ethnic salience and present these results in Table 2.9. The table also presents these results across the years, which show interesting patterns and p-values corresponding to testing the significance of the treatment

[10] I treat Kampala Metropolitan area, which includes Kampala and Wakiso as the one district as these two share the cities and there are clear boarders between these.

variables for arts majors: coethnic share ($\phi_1+\varphi_1$) and high-ability share ($\phi_2+\varphi_2$) in equation (2.6).

First, students with assumed high ethnic salience perform significantly lower than those with low assumed ethnic salience. However, this negative difference reduces over time and is no longer significant in the third year. That is, students of assumed high ethnic salience perform 0.63pp, significant at 1% lower than those of assumed low-ethnic salience in the first year, but the coefficient of this dummy increases to -0.17 and is no longer significant in the third year.

Second, the effect of coethnic share on students with low assumed ethnic salience is not significant across the years. However, Table 2.9 shows that students of high ethnic salience type benefit from a high share of coethnic peers in the first and second year. The differential effects of coethnic in the first, second, and third years are 2.088pp (significant), 1.323pp (significant), and 0.515 (insignificant), respectively. The table also reports the p-values of treatment variables at the bottom, which show that coethnic peers are important for students with high ethnic salience in the first year and second year. These results imply that adding five coethnic peers to a group of 25 increases academic performance by 0.42pp more for students assumed to be of high-ethnic salience type than those assumed to be of low ethnic salience in the first year. That is, adding five coethnic peers to a group of 25 leaders leads to a 0.44pp increase in academic performance, which is equivalent to a 0.05 standard deviation in the first year and is 2.5 times the average effect reported in Table 2.3.

Third, although positive, the differential effect of high-ability share by ethnic salience is small and insignificant. The share of high-ability peers has a positive and significant effect on students assumed to be of low ethnic and high ethnic salience in the first year, which persists and increases for both types of students in the third year.

These results indicate that students who might suffer from cultural and diversity shock when they arrive at MUK to study benefit more from coethnic peers than high-ability peers. Nevertheless, the effect of coethnic share decreases from the first to the second and disappears by the time the student graduates.

2.6.4 Robustness checks

One area of concern revolves around the potential impact of measurement error on estimates of the coethnic share in a peer group. As aforementioned, I use the probable coethnic share in a peer

Table 2.9: Differential Effect by Ethnic Salience: Coethnic vs High-ability Share

	Year One	Year Two	Year Three
Coethnic share	0.132	0.613	0.285
	(0.51)	(0.52)	(0.49)
High-ability share	0.716**	0.737**	1.044***
	(0.30)	(0.31)	(0.31)
High ethnic salience	-0.626***	-0.463***	-0.173
	(0.17)	(0.17)	(0.17)
Coethnics Share × High ethnic salience	2.088***	1.323***	0.515
	(0.48)	(0.48)	(0.46)
High-abillity Share × High ethnic salience	0.070	0.297	0.162
	(0.30)	(0.33)	(0.31)
p-val Coethnic share: High ethnic salience	0.000	0.000	0.130
p-val High-ability share: High ethnic salience	0.024	0.002	0.000
R-squared	0.328	0.384	0.380
N	321,375	343,761	330,158
Dorm FE	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes
Group controls	Yes	Yes	Yes

Notes: Data are from MUK and are restricted to students admitted to non-extension day majors from six colleges for 2009-2017, excluding 2010. A peer group comprises students admitted to majors within the same school and assigned to the same dorm. Each column is an independent regression, but the outcome is course grades in all regressions. All regressions control for own ethnicity, ability, and major, HS subject combination, and classroom FE. Individual controls include age, religious indicators, and graduating from the district of origin. Group controls include the leave-me-out averages of individual controls in addition to peer group size. SEs are parentheses and are clustered at the peer group level. The table also reports p-values for coethnic share and high-ability share of students assumed to be of high-ethnic salience. These tests correspond to $\phi_1 + \varphi_1$ and $\phi_2 + \varphi_2$ in equation (2.6).
* p<0.10, ** p<0.05, *** p<0.01

group to account for this. Nevertheless, I get practically similar results when I re-estimate the results using a single ethnicity corresponding to the category the algorithm is most confident about. I discuss robustness in relation to using student-level aggregated data and, thus, a different set of fixed effects in this section.

GPA as the Dependent Variable

I re-estimate the average effects at the student level, using GPA as the outcome (not course-level grades), and report these results in Table 2.10. Panel A compares the ethnic and high-ability shares within a student's peer group. In contrast, Panel B compares the effect of higher-ability coethnic peers to high-ability noncoethnic peers relative to low-ability noncoethnics. In addition, (1) is the same as the main effects regressions reported in Section 2.6.1 and is included for comparison purposes. The results obtained using GPA as the outcome are similar to those reported in Section 2.6.1. Naturally, the magnitudes of coefficients are different since the outcome variables are different. The results are consistent when I use school-by-year in place of classroom FE.

From Panel A, the effect of high-ability and coethnic share is positive and significant when the outcome is GPA Similarly, from Panel B, high-ability coethnic and noncoethnic peers positively and significantly affect academic performance. However, panel B shows the effect of low-ability coethnic is precisely estimated when I use GPA. As in Table 2.4 of the main effects, using GPA as an outcome also suggests that high-ability coethnic matter as much as high-ability noncoethnic peers although. Although the effect of high-ability coethnic peers is larger than that of high-ability noncoethnic peers in panel B column (2), the difference of 0.02 is not significantly different from zero.

2.6.5 Discussion and Contextualizing Results

I find that the share of high-ability and coethnic peers positively and directly affects academic performance, although the effect of the latter does not persist. The results reported in Table 2.3 suggest a mean peer effect size of 0.02 SD for both peer types. These are reduced-form effects based on dorm assignment, which is likely to be an underestimate of the true peer effect (treatment effect on the treated). This effect is comparable to what zimmer2003 finds at Williams College as Figure 2.5 shows. The effect Garlick (2018) finds at the University of Cape Town (UCT) using randomly assigned residential peers assignment is larger than what find, although his reported confidence intervals are large. The estimate is Garlick (2018) also reduced form effect because the author observes dorm assignments but not roommates. However, compliance is high in Garlick (2018) and is probably characterized by students who enroll at UCT from out of the city.[11]

Interestingly, I find strong coethnic reduced-form peer effects—equal to 0.05 standard deviations, especially for students of assumed high ethnic salience, which is comparable to the average

[11]The author mentions that people who do not live on campus in private residences, most likely with families (page 348).

Table 2.10: Coethnic vs High-ability Share: Outcome as GPA

	Outcome variable: % course grades	Outcome variable: GPA
Panel A: Coethnic vs High-ability		
Coethnic share	0.936**	0.112**
	(0.47)	(0.05)
High-ability share	0.735***	0.066**
	(0.28)	(0.03)
R-squared	0.328	0.277
N	321,452	25,298
Panel B: High-ability (Coethnic vs Noncoethnic)		
High-ability coethnic share	0.867***	0.091*
	(0.51)	(0.05)
Low-ability coethnic share	0.639	0.076*
	(0.39)	(0.04)
High-ability Noncoethnic Share	0.879***	0.081**
	(0.32)	(0.03)
R-squared	0.328	0.277
N	321,375	25,298
Classroom FE	Yes	No
School-by-year FE	No	Yes
Dorm FE	Yes	Yes
Individual controls	Yes	Yes
Group controls	Yes	Yes

Notes: Data are from MUK and are restricted to students admitted to non-extension day majors from six colleges for 2009-2017, excluding 2010. A peer group comprises students admitted to majors within the same school and assigned to the same dorm. Each column is an independent regression, but the outcome is course grades in all regressions. All regressions control for own ethnicity and dorm, major, and HS subject combination FE in addition to individual and group controls. Individual controls include age, religious indicators, and graduating from the district of origin. Group controls include the leave-me-out averages of individual controls in addition to peer group size. SEs are parentheses and are clustered at the peer group level.
* p<0.10, ** p<0.05, *** p<0.01

effect in Carrell, Fullerton, and West (2009) at the Air Force Academy. In short, in this setting with high ethnic diversity, I still find both ability and coethnic peer effects where high ethnic diversity is expected to dampen peer effects of higher ability. Ethnic diversity effects are unlikely to play a role at MUK.

While channels behind peer effects literature, in general, are unclear, I hypothesize on the mechanisms behind these results by discussing explanations for these results in this section based

Figure 2.5: Comparison to past Papers

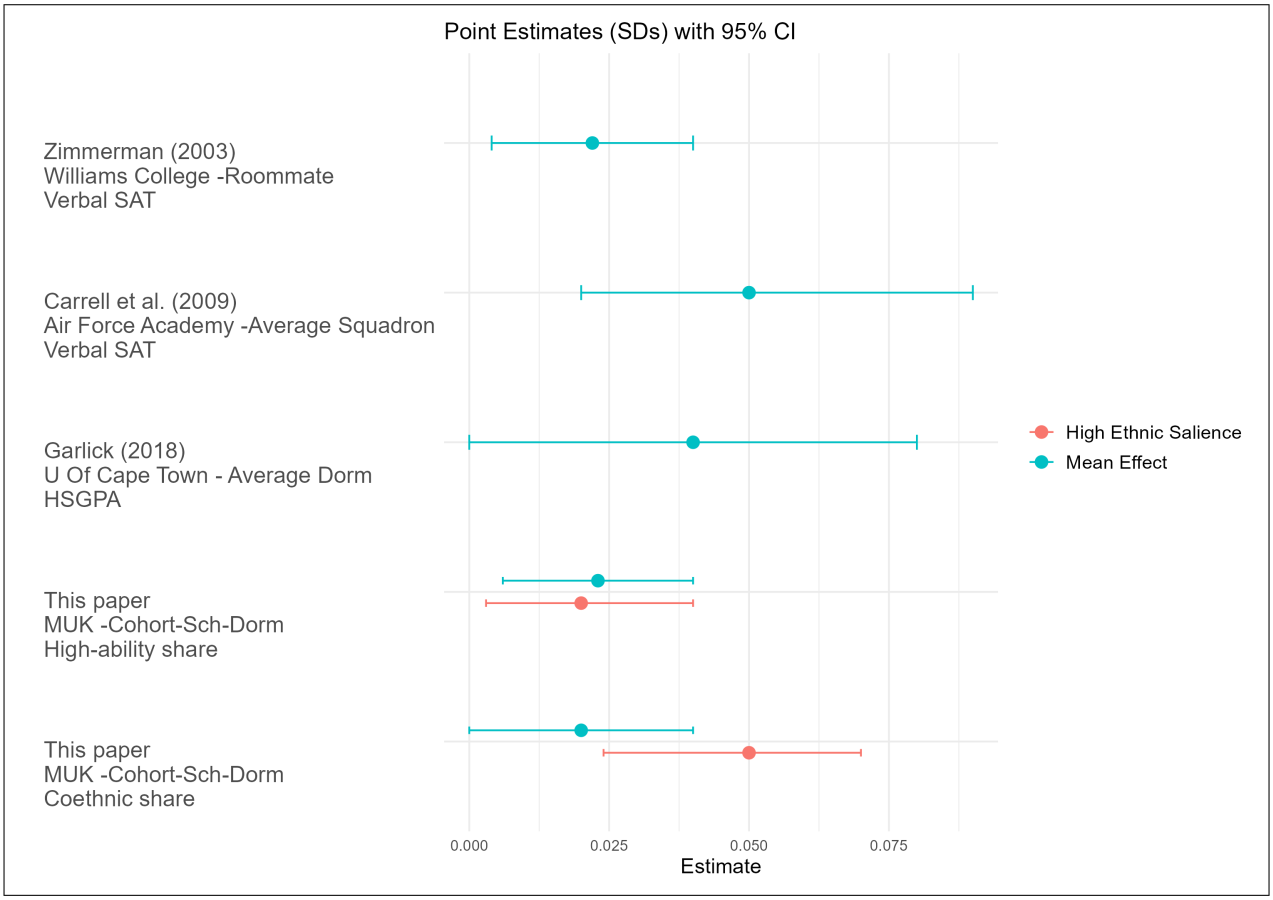

This figure compares my average estimate and estimate of students of assumed high ethnic salience to past papers with randomly assigned peers and significant average effects. Carrell, Fullerton, and West (2009) Table 3 column (6) page 452 reports a coefficient of 0.382 on peer SAT verbal, equivalent to a 0.05 increase in GPA. Additionally, Carrell, Fullerton, and West (2009) estimate is about 2.5 times that reported in Zimmerman (2003) Table 3 column ("First semester") page 17. Garlick (2018) Table 4 column (1) reports a coefficient of 0.216 on the dormitory mean high school GPA, equivalent to a 0.04 SD. Lastly, the effects of coethnic and high-ability share in Table 2.3 are about 0.02 SD increase in academic performance. Additionally, I report the peer effects on students of high ethnic salience type, showing that coethnic peer effects are about 2.5 times the average effect for these types of students.

on the mean and heterogeneous effects reported in the results section and the characteristics of this context.[12] The suggestive channels at play in this context that I discuss in this section include peer-to-peer learning, friendships, and psychological and cultural reasons.

Peer-to-Peer Learning and Study Behavior

Table 2.3 shows that one's own ability positively and significantly affects academic performance, an effect that persists into the final year of most majors. High-ability peers may influence the academic

[12] I do not test mechanisms directly. I am yet to start collecting primary data through surveys, for which I have already obtained IRB approval, including local IRB.

performance of their peers by facilitating peer-to-peer learning, such as leading discussion groups. This is especially important as office hours (professors or TAs) do not exist and because of the classical style of lectures in this setting. Students in need of extra help might rely on high-ability peers for additional assistance.

Students can identify their high-ability classmates through several methods, especially during the academic year advances. Firstly, student registration numbers differ by the enrollment scheme, such as merit scholarship status (high-ability). Secondly, it is common for newspapers to publish the names of the top students in the country (those with a high chance of obtaining a merit scholarship) once the national exams are out. However, this usually occurs several months before students enroll at university, and newspapers are not delivered outside the largest cities. Lastly, it is typical for students' course grades, especially in midterm marks, to be publicly posted on department noticeboards. Consequently, it becomes easy to identify and seek assistance from high-ability peers in ways that can impact academic outcomes.

Another potential explanation through which high-ability peers can influence others is by affecting study efforts. Several studies utilizing time-use data have examined how a student's study behavior is influenced by the study behaviors of their peers (e.g., Mehta, Stinebrickner, and Stinebrickner 2019; Frijters, Islam, and Pakrashi 2019). For instance, Mehta, Stinebrickner, and Stinebrickner (2019) show that students exhibit studious behavior if their peers, assigned randomly or connected through organic friendships, invest a lot of time studying at college or did during high school. It is conceivable that high-ability peers who have earned merit scholarships might have achieved it due to investing a significant amount of time into studying during high school or are doing so while at MUK. This intensified study behavior among high-ability peers could have a positive impact on the study behaviors of their peers.

Coethnic Friendships

Incoming freshmen can easily identify coethnic peers through physical features and cultural characteristics, including names and language. Shared ethnicity friendships are likely more organic due to shared identity since literature shows that coethnic bias exists in ethnically diverse societies. For example, Salmon-Letelier (2022) finds that ethnicity is important during friendship formation in Nigeria's state schools. Even studies outside the SSA report homophilous assortativity in student interactions in study groups and friendships based on gender and ethnicity (Jackson et al. 2022). Therefore, students might form ethnic-based friendships within randomly assigned groups explaining the suggestive evidence on why high-ability coethnic students might matter more for academic success as Table 2.4. This aligns with the Berea college freshman time-use (Mehta, Stinebrickner, and Stinebrickner 2019), which finds that using friends as peers is a stronger predictor of a student's propensity to study.

Nevertheless, the same table reports that high-ability noncoethnic peers also positively and significantly affect academic performance. Students may likely seek high-ability peers for academic help, irrespective of ethnicity. Thus, having high-ability coethnic peers is an added advantage because students may sort into friendships or study groups based on ethnicity when unaware of which of their peers are high-ability.

Cultural and Psychological Reasons

Lastly, these results also suggest cultural and psychological explanations at play. Many students migrating from rural districts might feel isolated as they navigate a diverse environment as they no longer belong to an ethnic majority. This could hamstring a sense of belonging for such students, which could have a negative effect on academic performance. As Table 2.9 shows, students of high ethnic salience perform lower than those of low ethnic salience in the first year.

In this case, a higher share of coethnic peers might be perceived equally or even more important compared to the share of high-ability peers by students experiencing a diversity shock. This mechanism might explain why I find coethnic peer effects in Table 2.3 are positive and significant, and even stronger in column one of Table 2.9 where the interaction of coethnic share and graduating HS from the district of origin is positive and significant.

If this mechanism is at play, this interaction should be even stronger for small groups (excluding Banyankore/Kiga and Baganda groups, which are the largest two groups that makeup 65% of the student population) as the smallest groups tend even to be more segregated as Figure 2.1 shows. Appendix Table A3 shows that the interaction is larger for smaller ethnic groups.

Nevertheless, this interaction could capture the influence of cultural shocks stemming from differences in diversity in the learning environment and between life in the city and rural areas. Beyond navigating diverse classrooms, migrated students encounter an urban lifestyle distinct from their rural upbringing. Additionally, it's conceivable that students who graduate from a high school within their district of origin might have predominantly resided at home, even though most Ugandan high schools offer boarding facilities. Such students might struggle to build a support network with peers, especially noncoethnic ones.

Table 2.9 also shows that the interaction's magnitude reduces from Year One to Year Three, and so does the mean effect of the coethnic share in Table 2.3. This pattern in the coefficients indicates that the importance of coethnic peers to students of assumed high ethnic salience goes away by the time they graduate. It is possible that cross-ethnic friendships emerge as these types of students acquaint themselves with peers through frequent interactions, making the coethnic share less important. The contact hypothesis, first introduced by William (1947), can explain this phenomenon.

Also, as students learn more about peers, ethnicity-based networks become less important compared to assortative matching based on attributes such as ability and study behaviors that matter more for academic success at college. This evolution of networks and information gain might explain why the effect of ability share increases with time.

2.7 Conclusion

Ethnic diversity has widespread and measurable impacts on a host of social, political and economic outcomes. In Sub-Saharan Africa, latent ethnic tension can deteriorate social trust and reinforce high coethnic favoritism. In the context of higher education, which brings students into close contact with ethnic diversity – often for the first time – ethnic heterogeneity may hamstring student collaboration and undermine academic performance with long-run implications. This paper provides causal estimates of peer effects on performance in the unique setting of higher education in Uganda, one of the region's most ethnically diverse and segregated countries.

I define a student's peer group as students admitted to majors within the same school in the same year who are assigned to the same dorm. This allows me to study the effects of peers with whom a student is likely to interact during school and non-school activities. Dorm assignments are random conditional on gender after a student is admitted, and courses are pre-determined at the time of admission before a student enrolls, providing an exogenous variation across peer groups.

I find that coethnic peers (irrespective of ability) and high-ability peers (irrespective of ethnicity) have a positive and significant effect on grades in the first year. However, the mean effect of coethnic peers does not persist until a student graduates.

These mean results mask significant heterogeneity in coethnic peer effects. First, I find strong and positive coethnic peer effects for students of high ethnic salience that do not persist until a student graduates. These are students who graduated from secondary schools in their districts of birth and have relatively limited exposure to ethnicities different from their own prior to arrival at campus. I also find a strong positive coethnic peer effect for high-ability students, not low-ability students, that persists. This suggests that the benefits of coethnic peers can be reaped by those who have the capacity to succeed academically. The results also suggest coethnic peers have a larger positive impact on girls than boys.

These results have a number of implications for higher education policy and administration in Uganda and, perhaps, in comparable settings with high ethnic diversity. First, the positive impact of high-ability peers on academic performance underscores the importance of fostering an environment that encourages peer-to-peer learning. For example, universities could implement optimal peer group assignments where low-ability students are mentored by high-ability students. Second, the positive effect of coethnic peers in the initial years on students assumed to be of high ethnic salience suggests that there could be benefit of implementing programs that facilitate cross-cultural awareness, shared cultural events, and increase a sense of belonging. Given the existence of ethnic student organizations in this setting, which suggests a degree of homophily that shapes student life, it is natural for incoming students of high ethnic salience to benefit from coethnic connections and support.

These results also suggest there might be a short-term cost to ethnic integration policies. For example, if a university peer group assignment algorithm breaks any homophily on ethnicity and enforces cross-ethnic mixing, it might have a negative effect on students who benefit from a higher share of coethnic peers, especially those assumed to be of high ethnic salience.

This paper points to several promising questions for future research. I find that a higher share of high-ability coethnic and noncoethnic peers increases a student's academic performance. At first glance, these findings suggest that college students at MUK may portray less coethnic bias during classroom interactions, such as study group formations that have an effect on economic outcomes.

In such cases, the peer effects in this setting work through channels, such as study effort, as some studies using colleges in the West (e.g., Stinebrickner and Stinebrickner 2006) report. However, these findings do not preclude other channels, such as coethnic cooperation and inter-ethnic competition. For example, high-ability coethnic peers might affect academic performance through cooperation with peers of shared ethnicity, while noncoethnic high-ability might increase competition where students of different ethnicities compete to the extent that increases academic performance.

Additionally, I study the first order of ethnic diversity on academic performance by focusing on coethnicity within a peer group. This paper does not study higher-order effects, such as the ethnic composition of noncoethnic peers, which is open for future research. For example, there might be an optimal pairing, tripling, quadrupling, etc., of ethnicities that could be beneficial or detrimental to academic performance. This kind of question requires going beyond studying the effect of ethnic diversity that would regress a Herfindahl index computed from ethnic shares within a student peer group on academic performance.

Lastly, this paper investigates short-term high-ability and coethnic peer effects by focusing on academic outcomes and finds that high-ability peers (irrespective of ethnicity) affect academic performance. However, it is unclear if a similar pattern of findings exists in the long term. Students may strategically engage during classroom interactions in a way that does not extend beyond the classroom. For instance, students might strategically select into study groups with higher-ability peers irrespective of ethnicity when doing homework but select into coethnic friend groups when forming non-education social networks. Cross-ethnic mixing at university may not change interethnic attitudes or social networks post-graduation if this happens. I focus these questions on the additional work I have initiated using the same setting of this paper.

References

Alesina, A., A. Devleeschauwer, W. Easterly, S. Kurlat, and R. Wacziarg. 2003. "Fractionalization." Journal of Economic growth 8:155–194.

Alesina, A., and E. La Ferrara. 2000. "Participation in Heterogeneous Communities*." The Quarterly Journal of Economics 115:847–904.

Alesina, A., and E. Zhuravskaya. 2011. "Segregation and the Quality of Government in a Cross Section of Countries." American Economic Review 101:1872–1911.

Barone, G., and S. Mocetti. 2016. "Intergenerational mobility in the very long run: Florence 1427-2011." Temi di discussione (Economic working papers) No. 1060, Bank of Italy, Economic Research and International Relations Area, Apr.

Bayer, A., S.P. Bhanot, E.T. Bronchetti, and S.A. O'Connell. 2020. "Diagnosing the Learning Environment for Diverse Students in Introductory Economics: An Analysis of Relevance, Belonging, and Growth Mindsets." AEA Papers and Proceedings 110:294–98.

Bhusal, B., M. Callen, R. Gulzar, Saad;Pande, S.A. Prillaman, and D. Singhania. 2020. "Does Revolution Work? Evidence from Nepal's People's War." EGA Working Paper Series No. WPS-116. Center for Effective Global Action. Unversity Of California, Berkeley, pp. .

Bloom, H.S. 1984. "Accounting for No-Shows in Experimental Evaluation Designs." Evaluation Review 8:225–246.

Bramoullé, Y., H. Djebbari, and B. Fortin. 2009. "Identification of peer effects through social networks." Journal of Econometrics 150:41–55.

Brunello, G., M. De Paola, and V. Scoppa. 2010. "PEER EFFECTS IN HIGHER EDUCATION: DOES THE FIELD OF STUDY MATTER?" Economic Inquiry 48:621–634.

Carrell, S., R. Fullerton, and J. West. 2009. "Does Your Cohort Matter? Measuring Peer Effects in College Achievement." Journal of Labor Economics 27:439–464.

Carrell, S.E., M. Hoekstra, and E. Kuka. 2018. "The Long-Run Effects of Disruptive Peers." American Economic Review 108:3377–3415.

Carrell, S.E., and M.L. Hoekstra. 2010. "Externalities in the Classroom: How Children Exposed to Domestic Violence Affect Everyone's Kids." American Economic Journal: Applied Economics 2:211–28.

Carrell, S.E., F.V. Malmstrom, and J.E. West. 2008. "Peer Effects in Academic Cheating." Journal of Human Resources 43:173–207.

Cavnar, W.B., and J.M. Trenkle. 1994. "N-gram-based text categorization."

Clark, G., and N. Cummins. 2015. "Intergenerational Wealth Mobility in England, 1858–2012: Surnames and Social Mobility." The Economic Journal 125:61–85.

Corno, L., E.L. Ferrara, and J. Burns. 2019. "Interaction, stereotypes and performance. Evidence from South Africa." IFS Working Papers No. W19/03, Institute for Fiscal Studies, Jan.

De Giorgi, G., M. Pellizzari, and S. Redaelli. 2010. "Identification of Social Interactions through Partially Overlapping Peer Groups." American Economic Journal: Applied Economics 2:241–75.

de la Cadena, M. 1995. Women Are More Indian: Ethnicity and Gender in a Community near Cuzco. Duke University Press.

Depetris-Chauvin, E., and R. Durante. 2017. "One Team, One Nation: Football, Ethnic Identity, and Conflict in Africa." Documentos de Trabajo No. 492, Instituto de Economia. Pontificia Universidad Católica de Chile.

Easterly, W., and R. Levine. 1997. "Africa's Growth Tragedy: Policies and Ethnic Divisions." The Quarterly Journal of Economics 112:1203–1250.

Eifert, B., E. Miguel, and D.N. Posner. 2010. "Political Competition and Ethnic Identification in Africa." American Journal of Political Science 54:494–510.

Foster, G. 2006. "It's not your peers, and it's not your friends: Some progress toward understanding the educational peer effect mechanism." Journal of Public Economics 90:1455–1475.

Frijters, P., A. Islam, and D. Pakrashi. 2019. "Heterogeneity in peer effects in random dormitory assignment in a developing country." Journal of Economic Behavior & Organization 163:117–134.

Gallen, Y., and M. Wasserman. 2023. "Does information affect homophily?" Journal of Public Economics 222:104876.

Garlick, R. 2018. "Academic Peer Effects with Different Group Assignment Policies: Residential Tracking versus Random Assignment." American Economic Journal: Applied Economics 10:345–69.

Gisselquist, R.M., S. Leiderer, and M. Niño-Zarazúa. 2016. "Ethnic Heterogeneity and Public Goods Provision in Zambia: Evidence of a Subnational "Diversity Dividend"." World Development 78:308–323.

Green, E. 2008. "District Creation and Decentralisation in Uganda."

Habyarimana, J., M. Humphreys, D.N. Posner, and J.M. Weinstein. 2007. "Why does ethnic diversity undermine public goods provision?" American Political Science Review 101:709–725.

Hjort, J. 2014. " Ethnic Divisions and Production in Firms." The Quarterly Journal of Economics 129:1899–1946.

Hooghe, M. 2007. "Social Capital and Diversity Generalized Trust, Social Cohesion and Regimes of Diversity." Canadian Journal of Political Science / Revue canadienne de science politique 40:709–732.

Hoxby, C.M. 2002. "The Power of Peers: How Does the Makeup of a Classroom Influence Achievement? (Research)." Education Next 2:57.

—. 2000. "The Effects of Class Size on Student Achievement: New Evidence from Population Variation*." The Quarterly Journal of Economics 115:1239–1285.

Hudson, M.C., and C.L. Taylor. 1972. World handbook of political and social indicators. Yale

University Press New Haven, Conn.

Håkansson, P., and F. Sjöholm. 2007. "Who Do You Trust? Ethnicity and Trust in Bosnia and Herzegovina." Europe-Asia Studies 59:961–976.

Jackson, M.O., S.M. Nei, E. Snowberg, and L. Yariv. 2022. "The Dynamics of Networks and Homophily." Working Paper No. 30815, National Bureau of Economic Research, December.

Jones, K.S. 2004. "A statistical interpretation of term specificity and its application in retrieval." J. Documentation 60:493–502.

Luhn, H.P. 1957. "A Statistical Approach to Mechanized Encoding and Searching of Literary Information." IBM Journal of Research and Development 1:309–317.

Mamdani, M. 2001. When Victims Become Killers: Colonialism, Nativism, and the Genocide in Rwanda. Princeton University Press.

Manski, C.F. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." The Review of Economic Studies 60:531–542.

Mehta, N., R. Stinebrickner, and T. Stinebrickner. 2018. "Time-Use and Academic Peer Effects in College." Working Paper No. 25168, National Bureau of Economic Research, October.

—. 2019. "TIME-USE AND ACADEMIC PEER EFFECTS IN COLLEGE." Economic Inquiry 57:162–171.

Miguel, E. 2004. "Tribe or Nation? Nation Building and Public Goods in Kenya versus Tanzania." World Politics 56:327–362.

Monasterio, L. 2017. "Surnames and ancestry in Brazil." PLoS ONE 12.

Morris, H.F. 1966. "The Uganda Constitution, April 1966." Journal of African Law 10:112–117.

NCHE. 2018. "THE STATE OF HIGHER EDUCATION AND TRAINING IN UGANDA 2018/19." A report on higher education delivery and institutions, The National Council for Higher Education.

Okunogbe, O. 2018. Does Exposure to Other Ethnic Regions Promote National Integration?: Evidence from Nigeria. The World Bank.

Ricart-Huguet, J., and E.L. Paluck. 2023. "When the Sorting Hat Sorts Randomly: A Natural Experiment on Culture." Quarterly Journal of Political Science 18:39–73.

Sacerdote, B. 2011. "Chapter 4 - Peer Effects in Education: How Might They Work, How Big Are They and How Much Do We Know Thus Far?" Elsevier, vol. 3 of Handbook of the Economics of Education, pp. 249–277.

—. 2001. "Peer Effects with Random Assignment: Results for Dartmouth Roommates*." The Quarterly Journal of Economics 116:681–704.

Salmon-Letelier, M. 2022. "Friendship Patterns in Diverse Nigerian Unity Schools." Comparative Education Review 66:709–732.

Sen, M., and O. Wasow. 2016. "Race as a Bundle of Sticks: Designs that Estimate Effects of Seemingly Immutable Characteristics." Annual Review of Political Science 19:499–522.

Ssentongo, J.S. 2016. "'THE DISTRICT BELONGS TO THE SONS OF THE SOIL': DECENTRALISATION AND THE ENTRENCHMENT OF ETHNIC EXCLUSION IN UGANDA." Identity, Culture and Politics: An Afro-Asian Dialogue 17:60–96.

Stinebrickner, R., and T.R. Stinebrickner. 2006. "What can be learned about peer effects using college roommates? Evidence from new survey data and students from disadvantaged backgrounds." Journal of Public Economics 90:1435–1454.

Tajfel, H. 1982. "Social Psychology of Intergroup Relations." Annual Review of Psychology 33:1–39.

Tornberg, H. 2013. "Ethnic fragmentation and political Instability in post-colonial Uganda: understanding the Contribution of colonial rule to the plights of the Acholi people in Northern Uganda."

Turyahikayo, B. 1976. Transafrican Journal of History 5:194–200.

UBOS. 2006. "2002 UGANDA POPULATION AND HOUSING CENSUS." Analytical report, Uganda Bureau of Statistics, October.

Uganda Bureau of Statistics. 2016. "The National Population and Housing Census 2014 – Main Report." Working paper, Kampala, Uganda.

Walton, G.M., and G.L. Cohen. 2011. "A Brief Social-Belonging Intervention Improves Academic and Health Outcomes of Minority Students." Science 331:1447–1451.

William, J., Robin M. 1947. The Reduction of Intergroup Tensions: A Survey of Research on Problems of Ethnic, Racial, and Religious Group Relations. New York: Social Science Research Council.

Zimmerman, D.J. 2003. "Peer Effects in Academic Outcomes: Evidence from a Natural Experiment." The Review of Economics and Statistics 85:9–23.

Zárate, R.A. 2023. "Uncovering Peer Effects in Social and Academic Skills." American Economic Journal: Applied Economics 15:35–79.

2.8 Appendix

2.8.1 Data Appendix

This section provides details that are not highlighted in the main data Section 4 of this paper.

Linking Student Data

MUK stores data on students' applications, admissions, and results in separate databases and offices. There is no unique identifier that can link databases in some cases.

STEP I: Computing GPA. My data cleaning process starts from the results database. These data list courses and course units (for some), and exam scores in percentages by program, department, semester, and year of study. They also list the calendar year when the exam was taken. These data cover 2008-2017. However, to match the admissions sample, I restrict the results sample to 2009-2017 years. I convert the exam scores from the percentage scale to letter grades using the information on the back of the transcripts and available in the code book. I then compute GPAs by semester and year.

STEP II: Merging with admissions. Each admitted student has two unique identifiers: student number and registration number. I use the latter to merge results and admissions data. There is a 93.9% merge rate at this stage.

Step III: Determining cohorts. The admissions and graduation programs are coded differently in many cases. The undergrad (graduation) program may admit students through different cohorts (e.g., evening and day classes). Take the graduation program "Bachelor of Science in Computer Science", for example; it is coded as "BCSCS". However, BCSCS students may be admitted to through two cohorts: day classes("CSC") or evening classes ("CSE"). This distinction was necessary because the cohort forms one's peer. I use the university codebook to ensure the admission, enrollment, and graduation programs are consistent. Since I restrict the sample to day majors, CSC appears in my final sample, while CSE does not.

Step IV: Merging with the name data. After correcting obvious misspellings in the names, I merged these data with data that predicted ethnicities. Merging on names in the training data gives a merge rate of 98.6%. Merging features produced by ML classification is irrelevant since ethnic predictions can be made for every surname.

Lastly, I deleted all the 2010 observations because the hall assignment is unavailable for many students admitted through a private scheme. The university officials in the admissions office mentioned that there was a problem/data bleach with the information system in 2010, where the university lost a lot of records.

2.8.2 Ethnic and Geographic Boundaries

The Ugandan parliament's gate has engravings of symbols and names of 15 administrative units at the time of independence from Britain. The administrative units were federal states, districts, or Territories (The Constitution of Uganda, 1964). The federal states were historical kingdoms, which included Ankole, Buganda, Bunyoro, and Toro, and the territory of Busoga. The districts included Acholi, Bugisu, Bukedi, Karamoja, Kigezi, Lango, Madi, Sebei, Teso, and West Nile. Coincidentally, these kingdoms and districts' boundaries followed ethnic/tribe boundaries that existed before the British colonial government but were exacerbated by British colonists.

However, the colonial government introduced a notion of a district as an administrative unit, which initially was a way to group similar ethnicities in geographical proximity. Kingdoms were historically centralized and ethnically segregated, with a traditional king as a ruler. However, this was different for districts. Some districts, such as Sebei and Bugisu, were ethnically segregated but followed a different system of local political leadership, such as clans or chiefdoms. There were also districts (e.g., West Nile and Bukedi) that were a cluster of several, and sometimes unrelated, relatively small ethnic groups. For example, the West Nile comprised mostly Lugbara people but included smaller ethnic groups, such as the Alur and Kakwa.

President Obote abolished kingdoms in 1966 for political reasons and changed the status of federal states to districts, and split the formerly powerful federal state (kingdom) of Buganda into four districts (Morris 1966).[13] Since then, the number of districts has increased to 135 over the years, with the highest increase happening under the current government for reasons such as service delivery and ethnolinguistic conflict management, among others. Some studies report political reasons as the most prominent explanations for new district creation (Green 2008).

[13]District is the second-largest unit of administration after the federal government. The districts divide into counties. Counties divide into sub-counties. Sub-counties divide into parishes/villages, which divide further into cells/villages.

Most importantly, new districts are curved out of existing districts at the time of creation. It has been rare to create a district by carving out counties that initially belonged to two separate districts over the years. Interestingly, albeit unsurprising, new districts tend to be more segregated by ethnicity (Ssentongo 2016). For example, the population of Nebbi is 96.2% of Alur ethnicity, although it was carved out of the West Nile district in 1974, which mainly comprised the Lugbara people. The creation of new districts sometimes begins with smaller ethnicities wanting to break away from the majority ethnicity in the original bigger district for reasons such as autonomy and bringing resources closer to them. But also, the government will offer a county a district status for political support.

I can trace current administrative units to historical kingdoms using publicly available data on administrative units from the Ugandan Ministry of Local Government. I complement the public data with data from the 2014 census from UBOS. The Census data contain the population breakdown by ethnicity for each district, confirming ethnicity within each district. That is, the census reports the number of each 66 ethnicities that reside in each district (i.e., 136 X 66 observations). I compute the proportion of each ethnicity in a district and rank these proportions from the highest to the lowest.

The top-ranked ethnicity informs the ethnic region that the district belongs to. The average proportion of the top ethnicity by population is 0.737 (the median is 0.813), indicating high ethnic segregation within each district. These UBOS data help me confirm the historic ethnic regions and give the final ethnic and geographic boundaries. I then create ethnic clusters by combining both the current and historical administrative units to give final ethnic geographical borders. Using just the 1962 districts and kingdoms that were created by the British colonial government would give wrong borders as the colonial sometimes bundled together ethnic groups that did not have centralized governments, such as those found in the eastern parts of the country

Specifically, when retracing the ethnic borders, the ethnicity with the highest proportion in a district based on UBOS data combined with historical settlement patterns supersedes these geographic boundaries established by the British colonial government. Additionally, This study ignores the smallest ethnicities within each district. Take Abim district, for example, the population of Abim is 87% of Karimojong ethnicity and geographically belongs to the Karimojong subregion. Using both UBOS data and historic settlement, this study identifies Abim within Karimojong borders when running the ML algorithm. However, Abim comprises other small minority groups, such as Gimara (0.033%). By ignoring ethnic groups that make up 13% of Abim's population, I am implicitly assuming that the smallest ethnicities are forced to assimilate with the largest ethnic groups within that district, or they are immigrant groups.

I use two formulas when allocating each district to the ethnic border (I) proportion of the highest ethnicity in the district and (II) ethnic fractionalization index. The two methods should give very similar borders. I use both for consistency. The ethnic fractionalization index introduced in Hudson and Taylor (1972) gives the probability that two randomly from a region (a district in this setting) belong to two different ethnic groups. I.e.,

$$FRAC_j = \sum_{e=1}^{E} \pi_{je} \left(1 - \pi_{je}\right), \tag{A1}$$

where j indexes a district, π_{je} is the proportion of ethnic group e in district j. Using UBOS ethnicity breakdown data by district, county, and sub-county, (I) and (II) are highly correlated (-0.981).

Table A1: Ethnic fractionalization in a district

	N	mean	sd
Ethnic fractionalization index	135	0.388	0.26
Max proportion in a district	135	0.727	0.22

From Table A1, the average proportion of the largest ethnicity in a district is 0.727, and the median is even higher (median=0.802). This implies that it is rare to find districts with equal shares of ethnicities. The average probability that two individuals are randomly selected from a district is low, and the median is also lower (0.345).[14] However, I compute this probability for the whole country, and I get 0.933. This is the same value reported in Alesina et al. (2003). Therefore, although Uganda is ethnically diverse as a whole, its subnational units are not. When constructing the training sample, I restrict districts where the ethnic fractionalization index is low (< 0.5), and the max proportion in a district is 0.7 and above.

Even though UBOS reports that Uganda has over 50 ethnic groups, 45 (68.2%) of the 66

[14] This is based on 66 ethnic groups in the census data. When I use ethnic clusters/language groups from Table A2, this index falls to 0.235

ethnic groups reported in 2014 census data contribute to less than 1% of the population each, and 22 ethnicities (33%) contribute a combined total of less than 1% of Uganda's population. The smallest ethnic groups are either non-Ugandan immigrant groups or indigenous groups. The immigrant ethnicities may be scattered across the country or segregated in the refugee resettlement areas.[15] The indigenous groups are tiny in that even though they are segregated, they only make up a small part of the district population. This leaves 32 unique ethnicities (out of 66) based on district and ethnicity clusters.

Students do not report ethnicity or places of origin during the application stage but their home districts. Although I observe home districts for most students, using reported districts would ignore cases of internal migration, especially rural-urban migration. Instead, I use students' surnames to predict their ethnicity as Ugandans' last names are almost usually in their native language, as Section 2.4.2 highlights. I combine ethnicities whose languages have high lexical similarity and mutual intelligibility to create a language group to proxy ethnicity.

Using language groups to proxy ethnicity has been used in several African studies to proxy ethnicity (e.g., Eifert, Miguel, and Posner 2010; Depetris-Chauvin and Durante 2017) as language and ethnicity usually overlap. The similarity in languages implies similarities in cultures, facilitating the ease of interaction in ethnically heterogeneous societies. Moreover, although not always, local languages in different follow a dialect continuum, which further informs my language groups/ethnicity. For example, historical and current Ankole and Kigezi people living in the SW part speak the same language but with different accents and are therefore combined to form the "Banyankore/kiga" ethnic group. Another basis for combining two or more ethnicities is historical. For example, the Tooro kingdom (Batoro) was historically part of the Bunyoro kingdom (Banyoro) until the early 19th century (Turyahikayo 1976). Therefore, Batoro and Banyoro form one ethnicity (language group). Combining groups that are mutually intelligible and similar reduces the ethnic groups to 16 groups. Another concern for the performance of the classification algorithm is how segregated ethnicities are. As Figure2.1 portrays, ethnicities within Uganda are geographically segregated.

[15]UNCHR ranks Uganda as the fifth largest refugee host nation. See this link: accessed 4/14/23

Table A2: Ethnicity/Language Group Composition

Ethnicity/language group	**Composition**	**Number**
Alur_Jonam	Alur, Jonam	2
SW	Banyankore, Bakiga, Bafumbira, Banyaruguru, Banyarwanda, Batagwenda, Barundi, Bahororo	8
Ganda	Baganda	1
Gisu	Bagisu and Babukus	2
Iteso	Iteso	1
Jopadhola	Jopadhola	1
Kakwa	Kakwa	1
Kelenjin	Pokot and Sabiny	2
Karimojong	Karamoja, Jie, Dodoth, Napore, Nyagia	5
Madi	Madi	1
Northern Luo	Acholi, Lango, Kumam, and Ethur	4
Nyoro	Batuku, Bunyoro, Batoro, Bagungu, Babwisi	5
Rwenzori	Bakonzo, Baamba	1
Samia_nyole_gwe	Banyole, Basmia, Bagwe	3
Soga	Basoga, Bagwere, Bakenyi	3
West Nile	Lugbara, Aringa	2
Extremely small	Vonoma(.008%), SoTopeth(.007%), Shana(.003%), Reli (.025%), Chope(.102%), Nube(.086%), Ngikutio(.017%), Mvuba(.009%), Mening(.008%), Lendu(.056%), Kuku(.140%), Kebuokebu(.161%), Bahehe(.012%), Gimar(.03%), Ikteuso(.041%), Batwa(.018%), Baruli(.565%), Banyabutumbi(.03%), Banyabindi(.049%), Aliba(.006%), Banyara(.142%), Nyangia(0.028%), Non-Ugandan(1.4%)	24
All		66

Notes: Source is the Uganda population and housing census of 2014. Groupings were informed using several sources as this section mentions.

2.8.3 Deriving the Reduced-Form Peer Effect

As described in the main text, this paper estimate the reduced-form peer effect based on random dorm assignment. In this section, I derive and discuss the relationship between this reduced-form estimate and the true underlying peer effect. Starting with equation (A1) and simplifying subscripts, we can write the individual specific effect of 'actual' high-ability share, $\widetilde{S}_i$ on student i's grade as

$$\text{(A2)} \qquad Y_i = \rho X_i + \phi \widetilde{S}_{iG} + e_i$$

where ϕ is the effect of the share of high-ability peers in a student's peer group on her academic performance. If I observed both random dorm assignment and actual (endogenous) dorm residence, it would be natural to use an IV approach to estimate the local average treatment effect of peers on academic performance, using dorm assignment to instrument for dorm residence as follows:

$$\text{(A3a)} \qquad \widetilde{S}_{iG} = \kappa_{10} X_i + \kappa_{11} S^H_{iG} + e_1 i$$

$$\text{(A3b)} \qquad Y_i = \kappa_{20} X_i + \kappa_{21} S^H_{iG} + e_2 i$$

where S^H_{iG} is the share of high-ability peers computed from peer groups as the result of the dorm assignment as in equation (A1) that may not be equal to $\widetilde{S}_{iG}$ because some students do not live in dorms. Equation (A3a) as the first stage capturing the effect S^H_{iG} on $\widetilde{S}_{iG}$, while κ_{211} captures the reduced form of the high-ability share due to dorm assignment. Substituting equation (A3a) into equations A2 will give:

$$\text{(A4a)} \qquad \kappa_{20} \equiv \rho + \kappa_{10}$$

$$\text{(A4b)} \qquad \kappa_{21} \equiv \phi \kappa_{11}$$

$$\text{(A4c)} \qquad e_{2i} \equiv \phi e_{1i} + e_i$$

Thus, the true high-ability peer effect (ϕ) is equal to $\frac{\kappa_{21}}{\kappa_{11}}$. That is, the IV estimate weights the reduced-form effect by the inverse of the first stage. Since I only observe dorm assignment, not residence, I am unable to recover this structural peer effects coefficient, so estimates captured in equation (A1) are reduced-form estimates of peer effects based on dorm assignment.

2.8.4 Additional Results

Differential Impacts by Ethnic Salience

The results presented in this section should be interpreted in conjunction with the effects in Section 2.6.3. I proxy high ethnic salience as graduating high school from one's district of origin. As illustrated in Figure 2.1, non-majority groups are even more segregated and might consequently encounter greater diversity shock when they relocate to the capital for university education. This is especially true since they are also the most underrepresented group at MUK. I present the differential effect by diversity shock in Table A3.

More on Robustness checks

Table A3: Differential effect by Ethnic Salience (Nonmajority) Coethnic vs High-ability Share

	Year One	Year Two	Year Three
Coethnic share	0.813*	1.037**	0.422
	(0.47)	(0.48)	(0.46)
High-ability share	0.605**	0.735**	1.000***
	(0.28)	(0.29)	(0.29)
High ethnic salience (nonmajority)	−1.394***	−1.043***	−0.644**
	(0.27)	(0.27)	(0.26)
Coethnic share × high ethnic salience (nonmajority)	3.254*	2.243	1.286
	(1.71)	(1.70)	(1.66)
High-ability share × high ethnic salience (nonmajority)	1.355**	1.095*	0.972*
	(0.56)	(0.59)	(0.52)
p-val Coethnic share (nonmajority): high ethnic salience	0.016	0.054	0.298
p-val High-ability share (nonmajority): high ethnic salience	0.026	0.084	0.173
R-squared	0.328	0.384	0.380
N	321,375	343,761	330,158
Dorm FE	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes
Group controls	Yes	Yes	Yes

Notes: Data are from MUK and are restricted to students admitted to non-extension day majors from six colleges for 2009-2017, excluding 2010. A peer group comprises students admitted to majors within the same school and assigned to the same dorm. Each column is an independent regression, but the outcome is course grades in all regressions. All regressions control for own ethnicity, own ability, and major, HS subject combination, and classroom FE. Individual controls include age, religious indicators, and graduating from the district of origin. Group controls include the leave-me-out averages of individual controls in addition to peer group size. SEs are parentheses and are clustered at the peer group level. Nonmajority ethnicities exclude the largest two groups (Banyankore/Kiga and Baganda). The table also reports p-values for coethnic share and high-ability share of arts degree. These tests correspond to $\phi_1 + \varphi_1$ and $\phi_2 + \varphi_2$ in equation (2.6). * p<0.10, ** p<0.05, *** p<0.01

Table A4: More Evidence against Selection

	Coethnic share	High-ability share
Age	0.000	0.000
	(0.00)	(0.00)
Anglican	-0.000	-0.002
	(0.00)	(0.00)
Catholic	0.001	-0.002
	(0.00)	(0.00)
Muslim	0.001	-0.003
	(0.00)	(0.00)
Seventh Day Adventist	-0.003	-0.004
	(0.00)	(0.01)
Pentecostal	-0.002	0.002
	(0.00)	(0.00)
High ethnic salience	0.001	0.002
	(0.00)	(0.00)
Other Religions	0.001	0.007
	(0.01)	(0.01)
High-ability	0.000	0.011***
	(0.00)	(0.00)
Peer group Size	0.000**	-0.000
	(0.00)	(0.00)
R-squared	0.420	0.263
N	25,323	25,323
Joint Fstat	0.84	2.15

Notes: Data are from MUK and are restricted to students admitted to non-extension day majors from six colleges for 2009-2017, excluding 2010. Each column is an independent regression that regresses either the coethnic share or high-ability share on all pre-university characteristics. All regressions include school-by-year FE (not classroom), ethnicity, and dorm FE. SEs clustered at the peer group level.
*p<0.1, **p<0.5, ***p<0.001

2.8.5 List of Figures

Figure 2.6: Distribution of Peer Group Sizes.

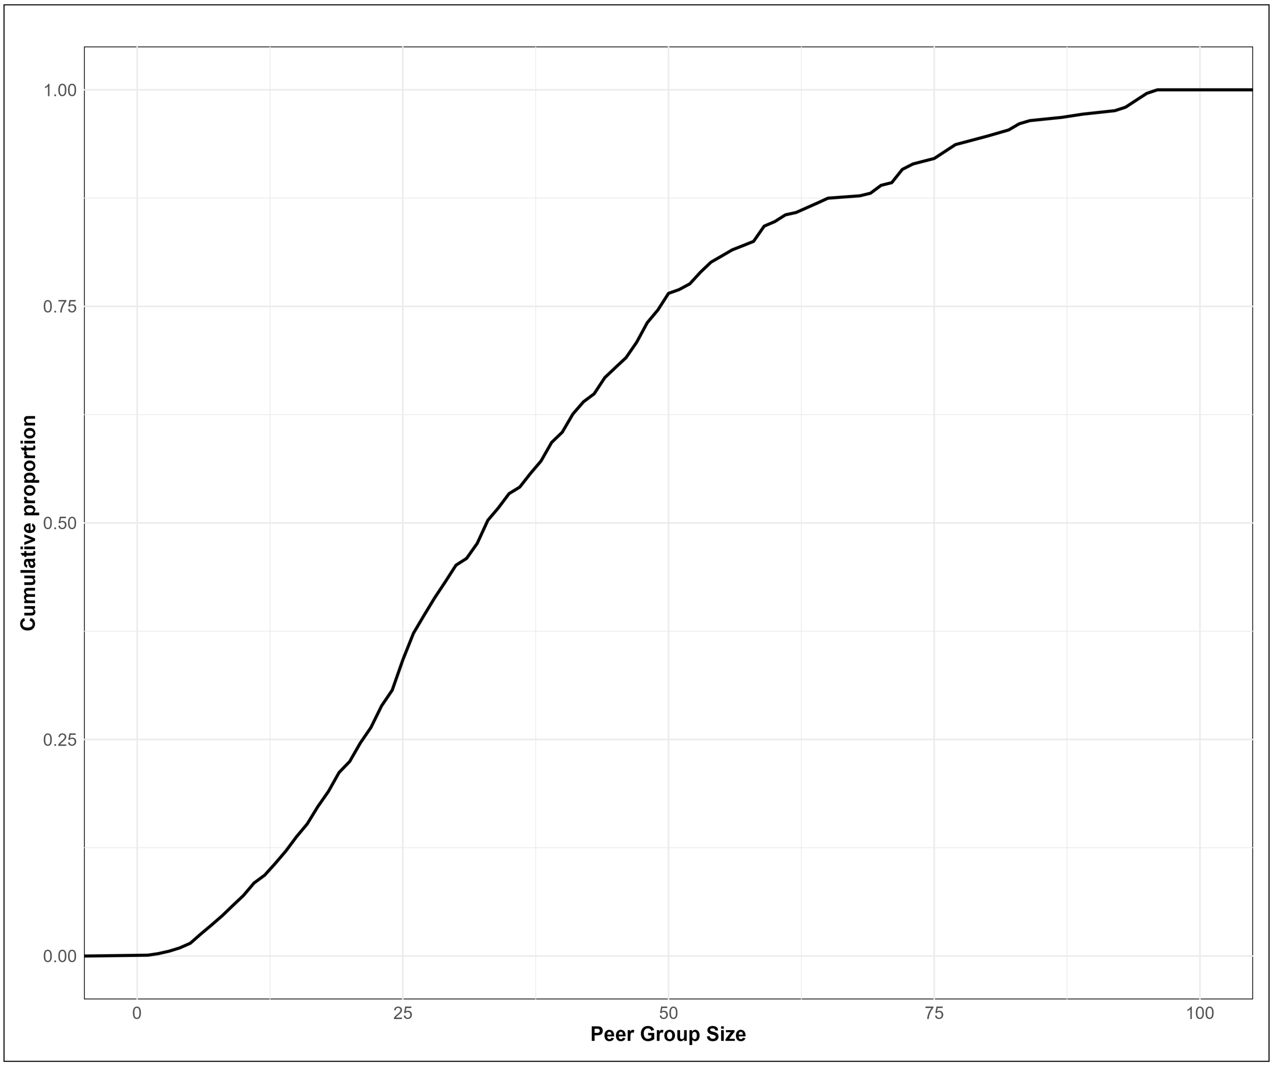

Notes: Data are from MUK and are restricted to students admitted to non-extension day majors from six colleges for 2009-2017, excluding 2010. A peer group includes students admitted to majors within the same school and assigned to the same dorm.

Chapter 3

Heterogeneity in Coethnic Peer Effects

3.1 Introduction

University administrators do not have control over endogenous social networks among students, although they can implement policies that directly influence students' peers through dorm assignments. For example, Makerere University's dean changed the dorm assignment system from alphabetical to random in the 1970s to avoid ethnic clustering within dorms (Ricart-Huguet and Paluck 2023). However, I find that a student's grades increase with a higher share of coethnic peers in Chapter One. Anecdotal evidence also shows that room assignments within each dorm are done to encourage interactions for academic success among students within the same major and school, although dorm assignments are random.[1]

Thus, an optimization problem for university administrators should seek to maximize the expected academic success and help promote national identity, not tribal identity, by the time students graduate and enter the job market. The latter is crucial in a setting like Uganda, which is characterized by high ethnic segregation, and university campuses are the only opportunity for students migrating from disparate regions to interact with peers of different ethnicities before they enter the workforce. Obtaining the optimal policy involves estimating the potential outcomes for different treatments and using these estimates to inform policy decisions (Athey and Wager 2021).

The average effect of coethnic peers reported in Essay One obscures important variations in

[1] My discussion with one of the dorm custodians at Makerere revealed that they intentionally assign rooms after receiving a list of students randomly allocated to each dorm. One criterion they consider is the proximity of students' majors and departments, both in courses and physical location.

interethnic interactions and other heterogeneity that may be crucial for policy. For example, the average effect does not explore any higher-order composition peer effects that may be detrimental or beneficial for student success. The Baganda and Banyankore communities, for instance, may be unfriendly toward each other, as these two ethnic groups have long competed politically.[2] It is not surprising that Baganda's presidential voting against the current regime has overwhelmingly increased in the last elections. Baganda may feel victimized because of persistent electoral losses. On the other hand, the Baganda may feel superior to northern minority groups. Thus, interethnic dynamics may differ for the Baganda if they are in a peer group with a large number of students of Banyankore descent compared to a group with a larger number of northern minorities.

Additionally, Essay One shows that coethnic peers are beneficial for academic performance, but a policy planner might be interested in the optimal number of coethnic peers that is most beneficial without hurting non-coethnic peers. Such a planner might be interested in learning about the non-linearities that might exist in the form of a dose-response function regarding the share or number of coethnic peers within a student's peer group.

To analyze the non-linearities and potential interethnic composition effects, I use casual forest estimation methods in Athey and Imbens (2016) and Athey and Wager (2019). Casual forest uses data-driven sample splits, reducing researcher bias in selecting the relevant heterogeneity dimensions. Additionally, the causal enable the capture of high-dimensional nonlinearities while avoiding overfitting by employing both training and estimation samples (the "honest approach").[3] Essentially, I estimate Conditional Average Treatment Effects (CATE) for each individual under this method by feeding the causal forest algorithm an estimation formula similar to my main estimation regression (equation 4) of Essay One.

My analysis reveals several findings. The predicted treatment effects are slightly nonlinear, with an increase in the coethnic share. The effect is negative and not significant when the coethnic share is small (less than 0.3) but positive and significant when the ethnic share is above 0.35. Additionally, the predicted treatment effects differ substantially by ethnic group. Interestingly, the

[2] In a recent mobilization tour in Buganda, the presidential candidate, Bobi Wine, seemed to subtly galvanize the Baganda to resist the Banyankore's occupation of the Baganda's ancestral land. The Banyankore have controlled Uganda's government since 1986, and Uganda's capital is geographically located in the Buganda kingdom. Consequently, the Banyankore have acquired extensive property on land belonging to the ancestors of most Baganda and built institutions to prolong their stay in political power, which could scar intergroup relations between these two groups.

[3] Moreover, treatment effect estimates using honesty fitting are asymptotically normal (Athey and Imbens 2016).

effect is largest for the largest ethnic group (Baganda), where 100% of their predicted individual treatment effects are to the right of the average treatment effect (ATE). Although positive, the predicted treatment effects are smallest for the second largest group, Banyankole, for which I may be powered to detect the treatment effects.

Additionally, the predicted treatment effects reveal significant gender differences across ethnicities. These differences are minimal among the Banyankore (the second-largest ethnic group) and most pronounced among the Basoga (the third-largest group). Notably, gender effects are not unidirectional: female Basoga students benefit more from the treatment than their male counterparts, while the reverse is true for the Baganda. When analyzing interethnic effects using a causal forest, the results do not show much variation in the predicted treatment or support the hypothesized interethnic effects, possibly due to the model's limited power in capturing heterogeneity related to interethnic composition.

In addition to the psychological channels discussed in Essay One, the heterogeneity analysis in this chapter reveals a high level of homophily. This effect is most pronounced among the largest ethnic group, which likely has the strongest ethnic ties, and is lowest among the ethnic group that controls the central government. Controlling the national government may cause the Banyankore to identify more with being Ugandan than with their ethnic group.

The rest of this chapter is organized as follows. Section 3.2 provides the empirical strategy, including explaining how I implement the causal forest algorithm. Section 3.3 gives the results, while I provide a discussion and conclusion in Section 3.4.

3.2 Empirical estimation

Since I am interested in analyzing the heterogeneity related to coethnic peer effects, I ignore high-ability peer effects for this section to reduce the dimensionality when estimating the CATE described in Section 2.1 below. when I run the specification (4) in Essay One without controlling for high-ability peer effects, coethnic peer effects do not change since ethnic and high-ability share are not correlated because of random assignment. That is, my CATE estimation hinges on the following equation.

$$y_{ijG} = \beta_0 + \phi_1 S_{iG}^{E} + \beta_2 \mathbf{X}_{iG} + \delta_j + \omega_f + \lambda_d + \gamma_s + \varepsilon_{ijG},$$

where y_{ijcG} is the GPA that student i of ethnicity j and belonging to peer group G obtained n the first year. Controls are similar to equation (4) of Essay One. S^E_{iG} is the probable coethnic share of i's peer group. The main estimation controls for δ_j, which is i's most probable ethnic group, $\mathbf{X}_{iG}$ is a vector of i's background characteristics and includes i's own ability. Additionally, ω_f, λ_d, and γ_s represent school-by-year, dorm, and high school subject combination fixed effects (FE). Lastly, ε_{ijG} is the error term. School refers to department or faculty as defined in Section 3.3 of Essay One.

The coefficients of interest are ϕ_1, which captures the effect of attending lectures and potentially living with coethnic peers in this setting. The identifying assumption is that conditional on ethnicity, school, gender, and cohort, the coethnic share is independent of unobservable and a student's characteristics. That is,

$$S^E_{iG} \perp (U, X) \mid \text{Ethnicity, School, Gender, Cohort}$$

I show why this identifying assumption is true in Appendix 3.5.1. For identifying S^E_{iG}, I only δ_j, ω_f and λ_d (since dorms are single-sex). I control for subject HS FEs to improve precision.

3.2.1 Estimating the CATE

As supervised machine learning techniques, casual forests predict heterogeneity in causal treatment effects to estimate CATE defined as, $\hat{\tau}_i = E[Y_{1i} - Y_{0i} \mid X_i = x]$, where Y_1 and Y_0 represent the potential outcomes for the i-th individual when treated and untreated, respectively, and X is a vector of observable characteristics. Causal forests as in the case of Athey and Imbens (2016) do not assume a specific functional form for the relationship between the outcome, treatment, and covariates, allowing for complex interactions and non-linearities. They can naturally handle heterogeneity in treatment effects, providing individual-specific estimates of the treatment effect, hence the subscript i.

The splitting criteria in causal trees are designed to maximize the difference in treatment effects across the resulting subgroups (leaves) while ensuring accurate estimation within each leaf. This involves choosing splits that lead to significant differences in treatment effects rather than just improving the fit of the outcome model, as in the case of random forests. Additionally, the algorithm estimates the individual treatment effects "honestly" and accurately. That is, in an "honest" causal forest, the data is split into two sets: one used for constructing the tree (finding splits) and the other for estimating the treatment effects within the leaves. This separation helps to provide unbiased estimates of treatment effects, as the estimation data was not used to determine the splits. As an advanced technique, honest causal forests ensure multiple trees are built, each time using a different subset of data for splitting and estimation, as mentioned. The results are then averaged to produce stable estimates. Honesty also eliminates any biasedness that would have resulted from overfitting.

I use the GRF R package by Athey, Tibshirani, and Wager (2019) to implement causal forests. I do not begin by fitting regression forests to estimate the nuisance functions for the conditional mean outcome and the treatment propensity score as in Athey and Wager (2019).[4] I am essentially estimating equation (1), which is a version of my main estimation equation in Essay One.

In this approach, I provide the outcome variable (Year One GPA), treatment variable (coethnic share), and covariates, such as gender, and include the FEs required for identification as Section Appendix 3.5.1 shows to the “causal forest" function, which internally handles the estimation of necessary nuisance parameters. The function then grows the causal forest by optimizing splits to maximize treatment effect heterogeneity while maintaining accurate predictions, as aforementioned. Additionally, I exploit GRF’s tuning parameters, such as setting the minimum node size to optimize the performance.

Lastly, since my treatment variable is continuous, the causal forest will provide a partial effect of coethnic share as in the case of Wooldridge (2010).[5] In my case, GRF non-linearly and non-parametrically uses a splitting criterion to maximize $\hat{\tau}_i = \mathbb{E}\left[\frac{\mathrm{Cov}(S^E, Y|X)}{\mathrm{Var}(S^E|X)}\right]$, where X is a vector of characteristics whose heterogeneity I am interested in and FEs.

Figure 3.1: Dose-Response Function

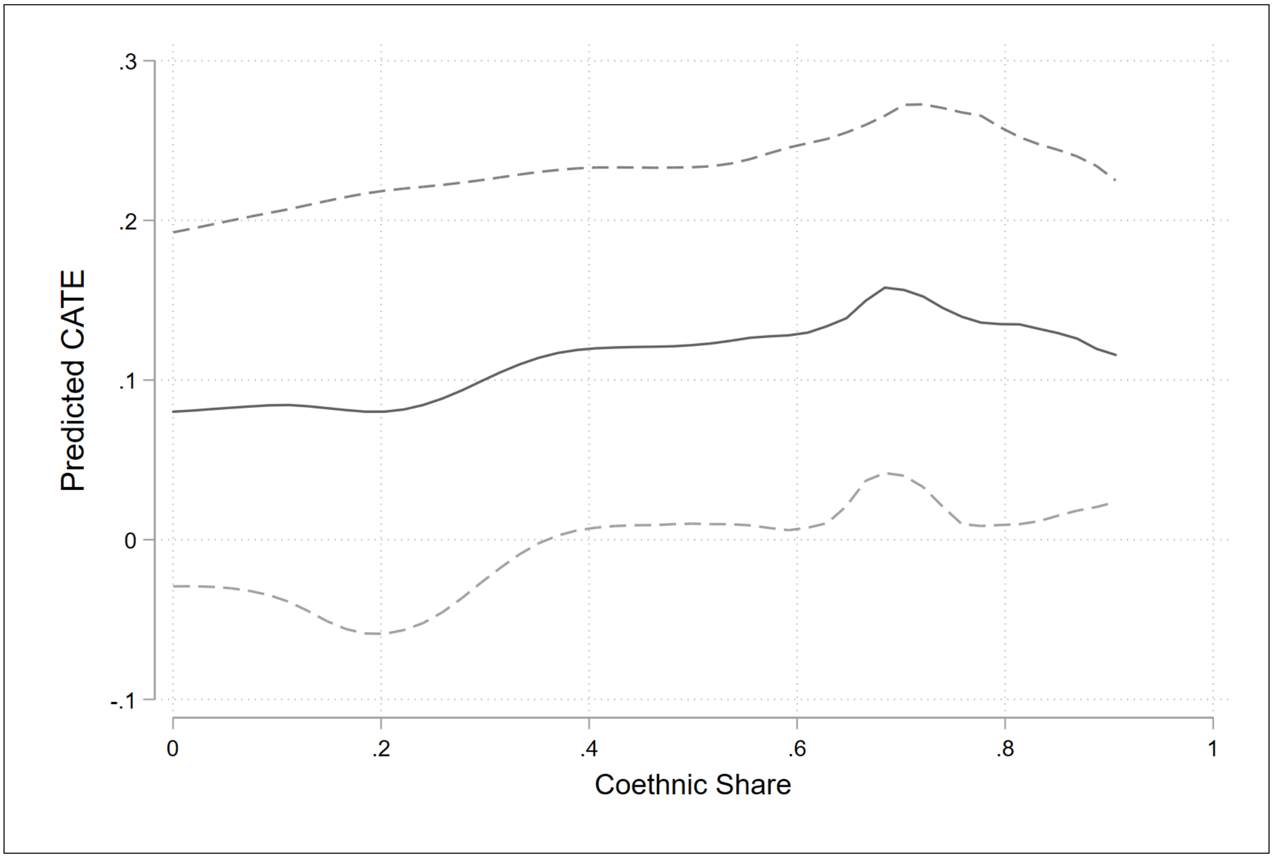

CATE is estimated using causal forest algorithms in the GRF package. The plot also includes the CI of the predicted CATE. Let $\hat{\tau}_i$ be individual i's predicted CATE. Each $\hat{\tau}_i$ is predicted with a variance (σ^2). Thus, the 90% CI is given by $\hat{\tau}_i$ +/-, qnorm (0.9) * sqrt(σ^2).

3.3 Results

3.3.1 Dose Response to Coethnic Share

From Figure (4) Essay One, we observe substantial variation in the treatment (coethnic share) by ethnic group. To visualize the relation between treatment and the predicted effects, I plot the individual predicted CATE against the treatment in Figure 3.1, including the 90% confidence interval (CI) using the variance predicted for each CATE using the casual forest algorithms.

I note several things. First, the CIs are large, which is expected—the authors of the causal forest report that the CIs tend to only converge in extensive samples. Second, most of the predicted

[4] The authors estimate those functions because the treatment units (schools) in their data exhibit selection. Essay One Table 2 provides evidence against selection. Although I use observational data, random dorm assignment ensures no selection.

[5] Wooldridge (2010), define the partial effect of a variable, w, $\mathbb{E}[y \mid w; X]$ as a derivative of $\mathbb{E}[y \mid w]$ with respect to w keeping X fixed

Figure 3.2: Difference by Ethnicity

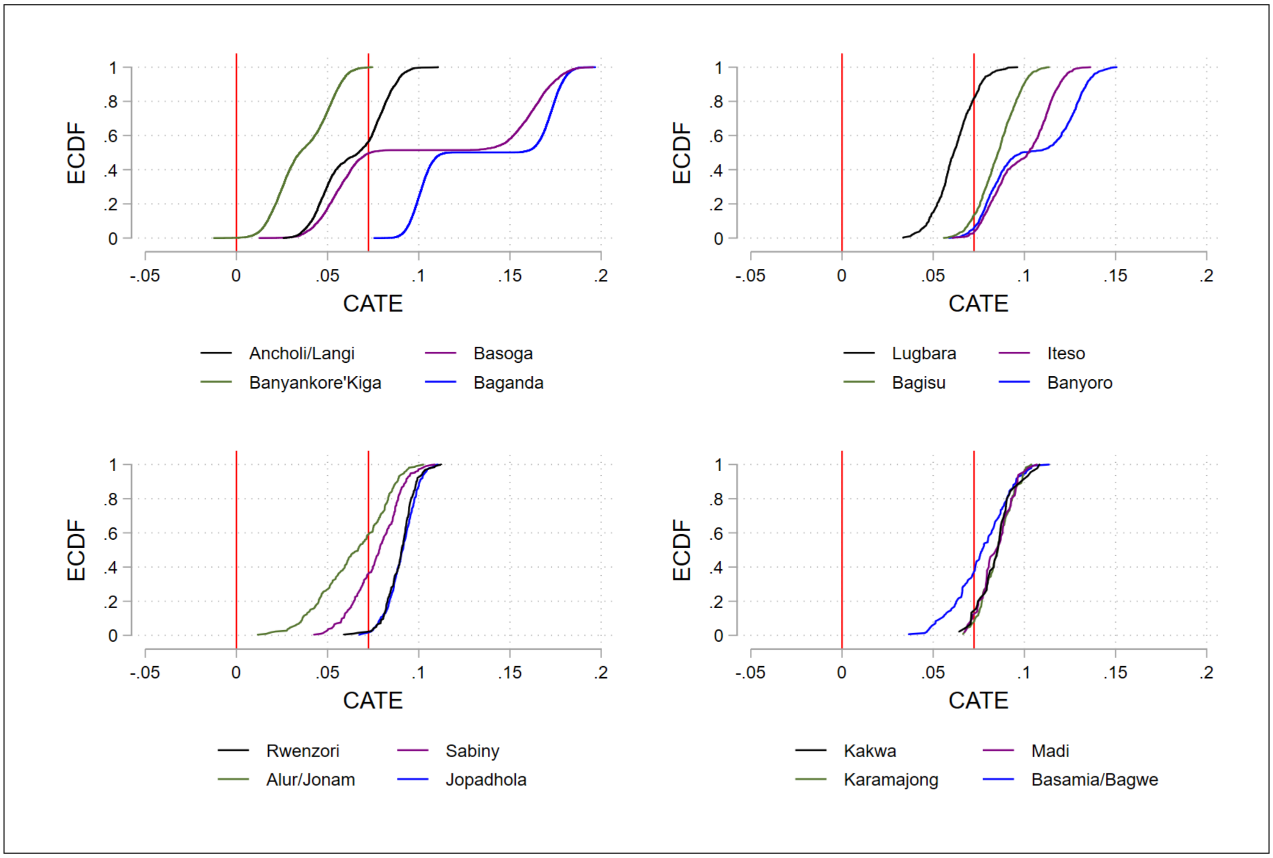

CATE is estimated using causal forest (cf) algorithms in the GRF package. The plot includes two vertical lines at zero and 0.0724. The latter corresponds to the average treatment effect from the causal forest: "average_treatment_effect(cf)"

treatment effects are positive. Third, the predicted CATE are increasing in the coethnic share. CATE is negative when the coethnic share is below 0.3, and positive and statistically significant coethnic share is around 0.35 or above. Lastly, the treatment effects seem to peak at a coethnic share of about 0.7, which corresponds to 18 coethnic peers in average-sized peer groups but the effect at 0.7 is not different from the effect at 0.6 coethnic share. When I break peer groups into small (size below the average) and large (size above the average), the partners are qualitatively similar to those portrayed in Figure 3.1.[6]

Figure 3.3: Differences by Ethnicity and Gender

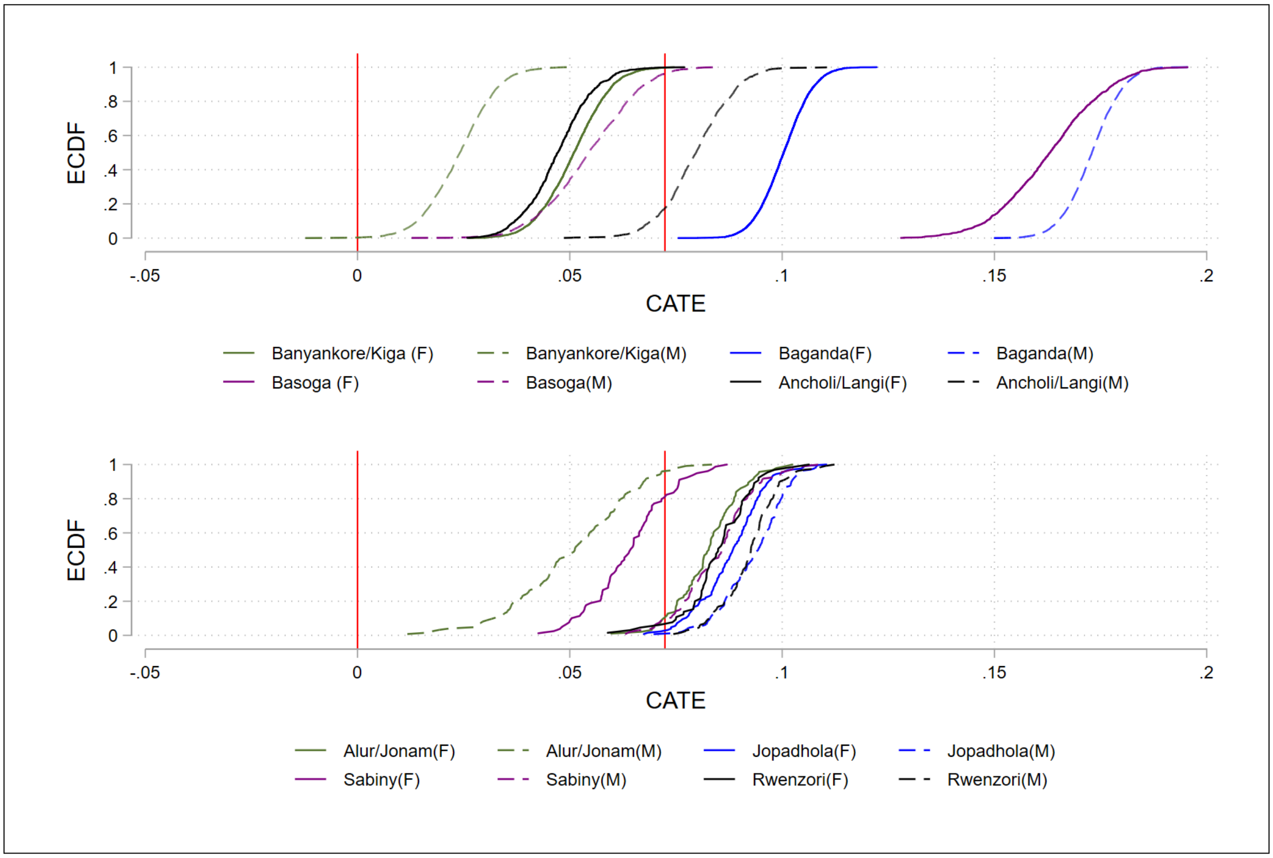

CATE is estimated using causal forest (cf) algorithms in the GRF package. The plot includes two vertical lines at zero and 0.0724. The latter corresponds to the average treatment effect from the causal forest: "average_treatment_effect(cf)". Within each panel, blue, green, purple, and black correspond to the first, second, third, and fourth ranking in terms of group sizes in my data. Additionally, the plot includes two vertical lines: one at zero and one at 0.0724, where the latter corresponds to the predicted average treatment effect.

3.3.2 Heterogeneity by Ethnicity

Figure 3.2 plots of the CDFs of the individual predicted treatment effects ($\hat{\tau}_i$) by ethnicity. I break the 16 groups into four panels: the largest four, the second largest four, and so forth. The causal forest also provides functionality for estimating the overall ATE, which is not just an average of the individual treatment effects. A corresponding coefficient plot is shown in Appendix Figure B5.

This figure shows substantial variation in the distribution of the predicted treatment effects especially. The top left panel shows that the Baganda have the largest treatment effect as the CDF of this group lies to the right of the ATE. Appendix Figure B5 shows that CATE for Baganda is significant. The second largest group, which we are part of, lies to the left of the ATE. The

[6] It is worth noting that the maximum coethnic share is .54 for the large groups, which is low as expected. The patterns are similar over this range of coethnic peers.

remaining two groups, Acholi and Basoga, show that 50% of the CATEs are below the ATE.

From the second panel, only 15% of the predicted effects are below the ATE for the Banyoro, Bagisu, and Iteso. In contrast, for Lugbara Baganda, 80% of the predicted average treatment effects lie to the left of the ATE, although positive for most individuals. We observe some variation in the third panel. However, the pattern is unclear, as the distributions of the smallest group, Rwenzori, and the largest group, Japadhola, overlap and stochastically dominate the other two groups.

Lastly, we do not observe variation across the smallest four groups, as the last panel portrays. However, this panel also shows that the predicted $\hat{\tau}_i$'s groups are large as the CDFs lie to the right of the ATE. We should also note that these groups are very underpowered to estimate the treatment effects accurately. For example, for the smallest group, Kagwa, we only have 50 instances of this ethnicity in the data, making it difficult to predict the treatment effects accurately.

3.3.3 Heterogeneity by Ethnicity and Gender

Gender norms may differ by ethnicity. For example, women in the central part of the country are culturally expected to be subservient to men, and some ethnicities in the eastern part of Uganda have cultural norms, such as genital mutilation, that differ by gender. I plot the CDF of the estimated treatment effects by gender and ethnicity in Figure 3.3. The top panel plots the largest four groups, while the bottom panel plots the third most prominent groups. I include only these groups for ease of comparison. Additionally, the ethnic groups in the bottom panel are very different from the ethnic groups in the top panel in terms of language, culture, and sometimes geography. I replicate a similar analysis in Appendix Table B1 for 14 out of sixteen ethnicities.

The top figure shows considerable differences in the estimated effects by gender across all ethnic groups. There is no overlap in the predicted effects by gender for each ethnicity. Also, the differences are not consistently larger for one gender across ethnic groups. For example, the predicted estimated effects for males are larger than the estimated effects for female students among the Baganda and Acholi. The opposite is true for Basoga and Banyankole. Lastly, the most prominent differences among Basoga are the smallest among the Banyankole. From the bottom panel, the treatment effects do seem to differ by gender across ethnic groups except Alur/Jonam. The bi-directional nature of differences by gender is the direction that requires further investigation I am unable to do right now due do data constraints as it may require qualitative data.

Figure 3.4: Interethnic Effects

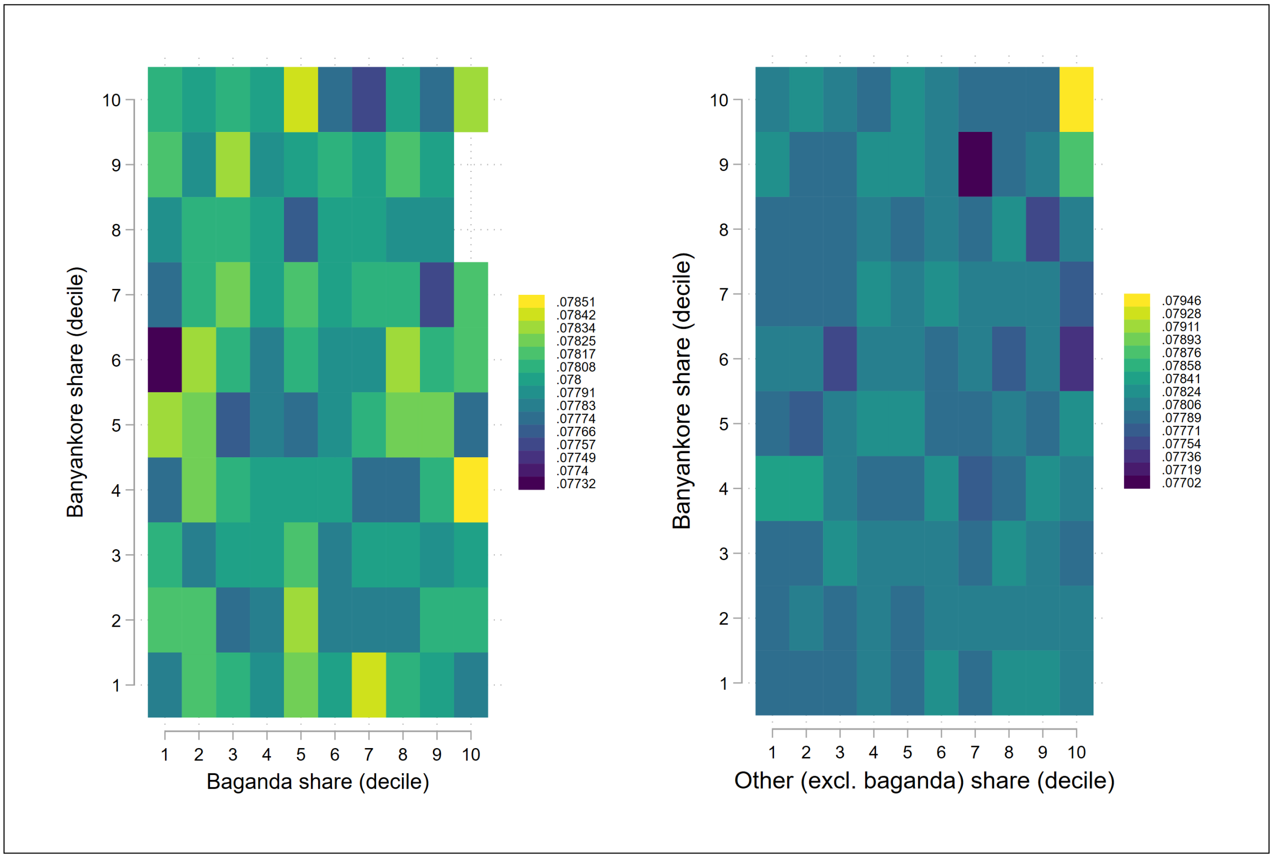

CATE is estimated using causal forest (cf) algorithms in the GRF package. This figure plots the mean predicted Conditional Average Treatment Effects (CATE) over the indicated pairs. The proportion of Banyankole and the proportion of Baganda in a peer group on the left and The proportion of Banyankole and the proportion of other ethnic groups (excluding Baganda) on the right. I estimate the correlations are equal to -0.44 and -0.258 in (I) and (II), respectively. Although computed correlations are somewhat low, they are significant. Thus, trees might split on one of these variables, especially in pair (I) when growing the casual forest.

3.3.4 Interethnic Effects

Figure 3.4 plots the predicted CATEs of two pairs share of Banyankore and Baganda and share of Banyankore and other ethnic groups. I hypothesize that if heterogeneity in the estimated treatment effects exists, it should show up in the left panel, not in the right panel. As mentioned, Baganda and Banyankole have a long history of political competition and may probably behave in a hostile way towards each other. I thus expected the predicted treatment effect should be high when the proportion of the other group is high, especially among the Baganda, as Section 3.3.2 showed that they benefit most from a higher share of coethnic pairs.[7]. Additionally, heterogeneity should not

[7]If the mechanism through with noncoethnic peers works is competition the expected predicted treatment could be high when the proportion of the other group is high

show up in the right panel since Banyankore and ethnic groups do not have a long history of animosity or political competition. Nevertheless, suppose other groups (mostly small groups) feel marginalized by the central government led by the Banyankore. In that case, a higher proportion of Banyankore might lead to a higher predicted treatment effect. However, the figure shows no significant heterogeneity in the predicted treatment effects. It is worth noting that the range in the predicted CATEs is low, as shown by the grid values in both panels. It is also worth noting that, given my sample size, I may be underpowered to grow to detect effects from trees grown for this prediction. When a parallel OLS equation with estimation, I also do not obtain significant differences.

3.4 Conclusion

University administrators aim to optimize assignment policies to boost student academic performance and help governments foster national identity over tribal identity, especially in ethnically segregated contexts like Uganda. This involves estimating outcomes for various treatments to inform policy decisions. My focus is on addressing coethnic peer effects and interethnic dynamics, which can significantly impact student success. For instance, historical political tensions between the Baganda and Banyankore groups may affect how students from these ethnic groups interact and their perceptions when they arrive on campus. Also, policymakers might be interested in nonlinearities in coethnic peers in a peer group and the potential discrimination or isolation faced by minority students. Ensuring the optimal assignment rule avoids harm and promotes positive interactions among diverse student groups is crucial.

I use causal forest to estimate expected treatment effects. This utilizes methods that leverage data-driven sample splits to minimize researcher bias in selecting relevant heterogeneity dimensions Athey and Imbens (2016); Athey and Wager (2019). This approach avoids overfitting by employing separate training and estimation samples. I estimate Conditional Average Treatment Effects (CATE) for each individual based on an individual's covariates.

My results reveal slightly nonlinear predicted treatment effects based on ethnic share, with insignificant effects when the coethnic share is below 0.3 but significant positive effects above 0.35. The largest effects are seen among the Baganda, while the Banyankole show the smallest positive effects, which portray high homophily.

The Buganda Kingdom is the strongest historical institution in Uganda, and it has survived until today. People from this group often show allegiance to their traditional king more than the central government. On the other hand, members of the Banyankore (the second largest group) might likely identify more as Ugandans than identifying primarily with their ethnic group due to their political status in the country. The people of Banyakore descent have controlled Uganda's political government for the last 38 years. Moreover, the Banyankore did not reinstate their historic kingdom as the Baganda did when the central government offered an opportunity to do so. It might explain why a large group, such as the Baganda, cares more about having coethnic peers than the Banyankore.

Lastly, the analysis also reveals Gender differences in CATE that vary by ethnicity, with notable disparities among the Basoga and minimal differences among the Banyankore. Interethnic effects do not show significant variations, possibly due to model limitations. The analysis highlights strong homophily, particularly among the Baganda, who maintain allegiance to their traditional kingdom, unlike the politically dominant Banyankore, who may identify more as Ugandan.

These results suggest that it might be feasible to optimize student performance by reallocating students across dorms, at least initially. This could involve leveraging nonlinearities in student characteristics to enhance academic outcomes. For coethnicity, segregating dorms might initially improve performance but could exacerbate long-term negative interethnic attitudes, contrary to the contact hypothesis highlighted in the first essay, which suggests that intergroup contact reduces prejudice. A thorough cost-benefit analysis, informed by the contact hypothesis literature, is essential to balance short-term academic gains with potential long-term social costs, ensuring that policies foster both academic success and social cohesion. Considering this cost-benefit analysis or trade-offs between academic performance and social cohesion is an area of future research.

References

Athey, S., and G. Imbens. 2016. "Recursive partitioning for heterogeneous causal effects." Proceedings of the National Academy of Sciences 113:7353–7360.

Athey, S., J. Tibshirani, and S. Wager. 2019. "Generalized random forests." The Annals of Statistics 47:1148 – 1178.

Athey, S., and S. Wager. 2019. “Estimating Treatment Effects With Causal Forests: An Application.” Observational Studies 5:37–51.

—. 2021. “Policy Learning With Observational Data.” Econometrica 89:133–161.

Ricart-Huguet, J., and E.L. Paluck. 2023. “When the Sorting Hat Sorts Randomly: A Natural Experiment on Culture.” Quarterly Journal of Political Science 18:39–73.

Wooldridge, J.M. 2010. Econometric Analysis of Cross Section and Panel Data, 2nd ed. MIT Press, Accessed: 2024-06-24.

3.5 Appendix

3.5.1 Deriving the Identifying Assumption

Let $Y = f(S^E, X, U)$ be the GPA production function, whereS^E denotes coethnic share, X denotes covariates (including ability A), U denotes unobserved factors. Also, let F denote school/faculty.

Conditional on gender and cohort:

1. Dorm assignment: $D \perp (U, X, F, E)$

   $$\Rightarrow D \perp (U, X, E) \mid F$$

   This is true because unconditional randomization implies conditional randomization. That is, dorm assignment is independent of unobservable, individual covariates, school, and ethnicity, which implies that conditional on school, dorm assignment is independent of U, X, and E

2. Peer group definition: $G = f(F, D)$

   $$\Rightarrow G \perp (U, X, E) \mid F$$

   In Essay One, I define a peer group comprised of students admitted to the same school and randomly assigned to the same dorm within each cohort and gender. Thus, given a random dorm assignment, my peer group is independent of U, X, and E conditional on school.

3. Peer measures:

   $$S^E = f(G, E)$$

   $$\Rightarrow S^E \perp (U, X) \mid F, E$$

   The peer measure, the coethnic share is a function of peer group and the number of coethnic peers. Combining (1.) and (2.) gives the identifying assumption in (3).

3.5.2 Additional Figures and Tables

Figure B5: Coefficient plot by Ethnicity: CATE

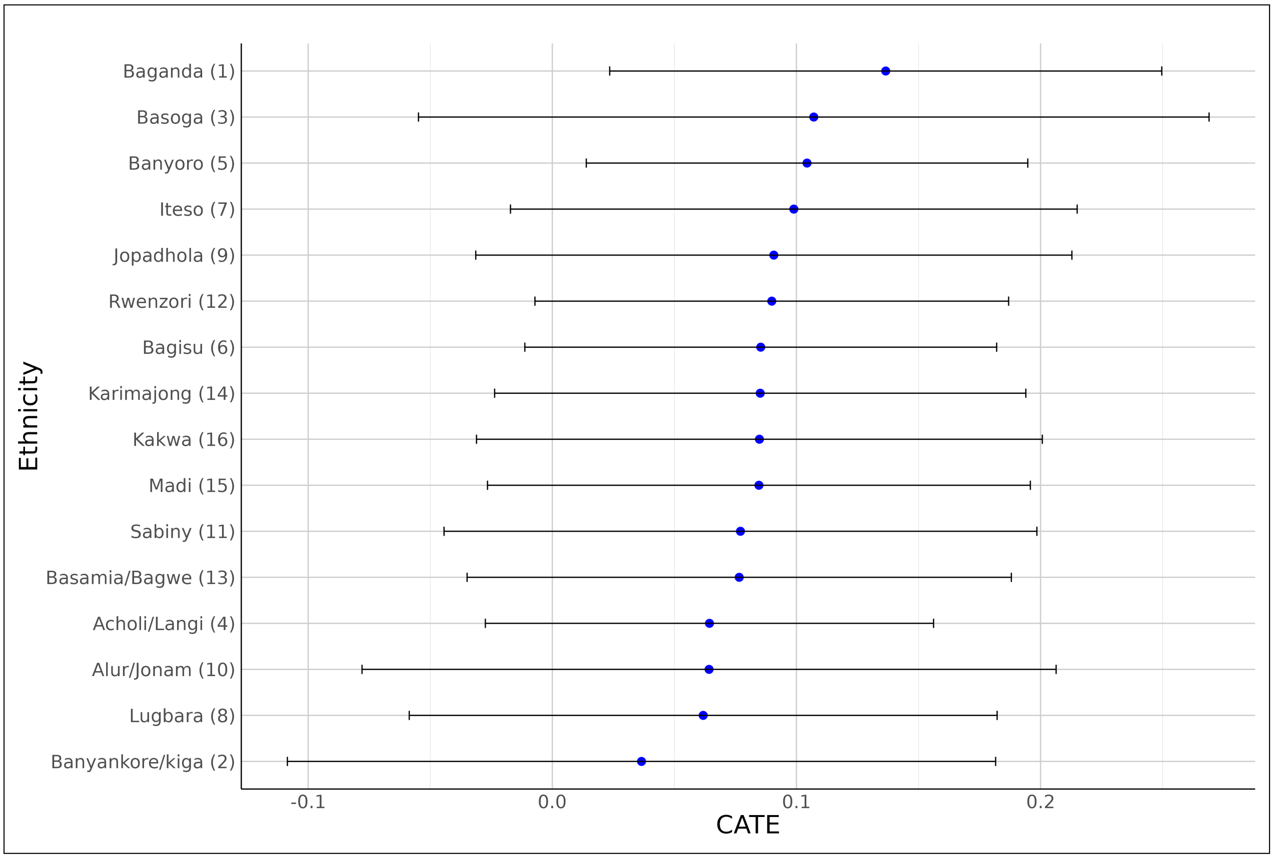

CATE is estimated using causal forest (cf) algorithms in the GRF package. This figure gives a coefficient plot by ethnicity along the 90% CI, computed as $\hat{\tau}_i$ +/- qnorm*sd(τ_i). Each $\hat{\tau}_i$ is predicted with its variance. Ethnicities are arranged in terms of the mean CATE from the lowest to the largest. The numbers in parentheses correspond rank of the ethnicity in terms of the sample size. These ranks should mirror the relative sizes in the country as Figure 2 of Chapter One. As aforementioned the CI are large, possibly driven by the high dimensional data and a not-so-large sample.

Table B1: Gender Differences CATE by Ethnicity

Ethnicity	**Female**	**Male**	**Difference**
Basoga (3)	0.163	0.054	0.109
Banyankore/kiga (2)	0.051	0.024	0.027
Alur/Jonam (10)	0.083	0.050	0.032
Jopadhola (9)	0.088	0.093	-0.005
Rwenzori (12)	0.085	0.092	-0.007
Karimajong (14)	0.079	0.091	-0.012
Lugbara (8)	0.054	0.067	-0.013
Bagisu (6)	0.077	0.093	-0.016
Sabiny (11)	0.064	0.085	-0.021
Basamia/Bagwe (13)	0.064	0.085	-0.021
Iteso (7)	0.082	0.112	-0.029
Acholi/Langi (4)	0.047	0.080	-0.032
Banyoro (5)	0.082	0.127	-0.046
Baganda (1)	0.100	0.173	-0.073

Notes: Table presents the AVERAGE CATE by ethnicity excluding the smallest two groups with a sample size below 100. The numbers in the parentheses are ranks of the ethnic group size. For example, Baganda (1) refers to Baganda as the largest ethnic group. I do not include the P-values testing the significance of the difference. Given the range in the predicted CATE is tight as shown by the range in Figure 3.2.

Chapter 4

Beliefs and the Demand for Employee Training

4.1 Introduction

The classic Becker (1962) model of human capital predicts that firms will not invest optimally, leading to the underprovision of general skills training. This suboptimal equilibrium is the result of the fact that, in perfectly competitive labor markets, firms will pay employees less than their marginal product in order to recoup the costs of training, which will induce employees to leave for other firms. Thus, government programs that subsidize training are common, particularly in rich and middle-income countries. Underprovision of human capital in firms, however, is particularly relevant in developing countries since firms in those settings are characterized by low productivity, which is a likely constraint to growth (Hsieh and Klenow 2009). Information frictions in such settings, however, may create unique distortions in labor markets. Thus, even if firm training is subsidized, firms may not choose the optimal employee to train.

In this study, we examine the sub-optimal provision of employee training using a novel experiment. We offer free training to firms belonging to one of the most critical skill-intensive manufacturing subsectors in Uganda, metal fabrication, and study how owners select workers for training. Specifically, we study if owners choose the socially optimal worker that would push the metal fabrication subsector's production possibility frontier outward the most or if they behave individually rationally by selecting a worker to maximize the firm's profits but is not necessarily the worker whose quality would not improve the most from training. Given that our training is free and designed to also minimize non-monetary costs, firm owners in our study sample should afford to pay a marginal product to workers post-training, ameliorating the anticipated friction of separation by trained workers.

We carry out data collection with metal fabrication firm owners, including incentive-compatible selection for training. We also elicit incentive-compatible owners' perceived quality of their workers at baseline and endline with or without training. Additionally, we collect perceived profitability from training each of their worker and data that proxies worker ties with the firm, such as if the worker is a relative, the owner's trust in the worker, and the perceived likelihood of separation. From the workers' side, we elicit incentive-compatible demand for training. Specifically, we ask workers to request an amount of money they are willing to accept (WTA) to attend training. We pay winners their WTA if randomly assigned to receive training when they attend training. Lastly, we test workers' quality using an objective measure (practical tests) scored by assessors from one of the prominent government vocational institutes.

We randomize small metal fabrication firms (4-14 employees) in our evaluation sample into treatment and control and offer a training program to the treatment group using a curriculum carefully designed through consultation with metal fabrication experts and lecturers from one of the prominent vocational training institutions in Uganda. Like most manufacturing firms in the developing world, metal fabrication firms in Uganda are small. Small firms are predominant in the Ugandan economy, where SMEs account for 90 percent of private-sector production according to Uganda's Bureau of Statistics (UBOS 2014). This is typical of most developing economies where SMEs contribute up to 40 percent of GDP, according to the World Bank. The productivity differential between SMEs and large firms – especially in industries with economies of scale – is particularly important in developing countries where small firms are also less likely to transform into large firms (Van Biesebroeck 2005; Olafsen and Cook 2016).

Our preliminary analysis reveals that, on average, owners believe that our training program has a positive benefit (measured by a perceived improvement in quality) on their workers. However, they are less likely to select a worker who would increase the productivity of the metal fabrication subsector the most. This is a worker whose skills would increase most at the firm. Instead, we find that owners are more likely to select the worker with the highest perceived profitability post-training but would not improve most from training.

That is, we find that owners select workers who they trust and who have strong ties to the firm. Specifically, owners rank a family member 10.1 percentiles higher for training relative to a non-family member. We also find that a one-point increase in perceived trust is associated with an 8.1 percentile higher rank and a similar increase in perceived risk of separation is associated with a 1.1 percentile lower rank in terms of worker selection for training. All these coefficients are significant at the one percent level. Third, we find that strong ties to the firm and trustworthiness are also significantly and positively correlated with perceived profitability from training a worker. Yet, there is a negative association between perceived profitability and perceived teachability.

Put together, our results reveal that even though our intervention provides a free training program, eliminating credit constraints on the side of the owner to provide training, and thus, owners could afford to pay the workers their post-training marginal profit, owners are not socially optimal. Instead, they are individually rational in that they select workers with the largest gap between post-training marginal product and the wage they can get away with paying without forcing that worker to leave the firm. Such employees typically have strong ties to the firm, such as relatives or workers who are highly perceived to be reliable by the owner in our setting. On the contrary, workers have a different objective function in that they demand training to maximize their marginal product and, consequently, their lifetime earnings. Our results show that the workers' demand for training does not strongly align with the owners' selection for training.

Our study contributes to several strands of literature. First, we contribute to studies documenting the effectiveness of several training programs in the developing world. Most of these papers study vocational programs.[1] For example, Alfonsi et al. (2020) compares firm-provided training to vocational training and finds both types of training improve employment and earnings outcomes for Uganda's "disadvantaged" youth starting in the labor market, but the impact of vocational training shows almost twice that of apprenticeships. We contribute to this literature by studying current firm employees and how anticipated frictions affect worker selection for training, and thus effective-

[1] For instance, McKenzie (2017) reviews different training programs across different settings, Card et al. (2011) studies a youth training program provided by the government in the Dominican Republic, Cho et al. (2013) study vocational training programs in Malawi, Hirshleifer et al. (2016) study effects of vocational training targeted to unemployed youth in Turkey on labor market outcomes. Most of these papers report modest treatment of the training programs they study.

ness of training programs in settings. We offer training to firms in one of the crucial subsectors in Uganda, and we find that owners do not choose the workers that would improve most from training and potentially have larger labor market outcomes. Instead, they select workers with low perceived gains who have strong ties to the firm.

We also contribute to the literature on the under-provision of training in skills that are transferable between firms (Acemoglu and Pischke 1998; Becker 1962; Prendergast 1993; Acemoglu 1997). We provide two possible explanations for these settings. First, our descriptive analysis shows that most employees entered the firm through apprenticeship, suggesting that owners have confidence in their ability to train low-quality hires. However, firm apprentices may not be effective for the industry because they lack standardization. Also, recent literature shows that training in similar settings increases the separation of trained workers (Brown et al. 2024; Frazer 2006). We provide experimental evidence that owners do not select workers with the highest gains from training even though they could potentially afford the marginal product in anticipation of separation.

Perhaps the more related, Cefala et al. (2023) study the under-provision of training in agricultural markets using two experiments. In the first experiment, both the control and treatment farmers receive incentives to train workers on their farms, but the treatment payouts are conditional on the farmer attending training. They find that tying the incentives to actually providing training reduces under-provision of training, but trained workers and non-training firms appropriate the returns to training. In the second experiment, they tie trained workers' incentive payouts to the worker working for a farmer who provides training. They find that this reduces under-provision of training even though owners do not receive financial incentives to train. The results can be interpreted that reducing the risk of separation and "poaching externality" increases training provision. Similarly, we find that owners prefer to train workers with stronger ties to the firm. Although the worker trained in their second experiment is chosen by the farmer, our experiment differs in that we observe the choice set over which the owner is deciding on who to send to training. For example, we observe if owners are deciding between a worker who is a relative and one who is not. That is, we can explain why training programs may not be effective because we observe the opportunity of training one worker over the other (we can compute the unrealized gains from owner choices). Lastly, our experiment goes beyond owners and studies worker decisions.

4.2 Context and Conceptual Framework

4.2.1 Metal Fabrication in Uganda

Micro, Small, and Medium Enterprises (MSMEs) are a strategic part of Uganda's policy for economic development. The National Development Plan II of Uganda's government (GoU) underscores the role of MSMEs in wealth and job creation and outlines government objectives to provide institutional support for MSME growth (GoU 2015a). According to the Uganda Bureau of Statistics (UBOS), MSMEs constitute over 90 percent of the private sector in Uganda and contribute 18 percent of the GDP. Additionally, manufacturing-related firms, including welding and steel, made up 8 percent of all MSMEs in 2011 and employed 15.4 percent of the workforce in 2012 (GoU 2018). Using statistics from UBOS, metal products (steel and fabrication) are among the top seven manufacturing sectors, according to the Index of Production. The same report shows that metal and steel fabrication grew substantially, over 10 percent, from 2011 to 2017, reflecting the dynamic nature and potential of the manufacturing industry in Uganda.

Despite the economic importance and potential for the economy, Ugandan MSMEs face a myriad of challenges. GoU (2015b) summarises Uganda's MSME policy and lists access to credit, lack of adequate technical skills, and informality as the top three challenges MSMEs face hindering their growth. More recently, and specifically for metal fabrication in Uganda, Bassi et al. (2023) show that another challenge that limits growth is the customer-tailored nature of Uganda's welding businesses that also stifles specialization.

During our listing exercise, we visited around 2,053 firms in Kampala and neighboring suburbs and found that less than 40 percent of firms employed at least four workers. In a different sample, Bassi et al. (2023) reports an average number of six employees in welding firms in Kampala even after oversampling the largest firms. Small welding firms in our setting are likely to face internal constraints long reported in the literature in other settings that hinder growth, such as sub-optimal management practices in small businesses across several developing countries (McKenzie and Woodruff 2016).[2]

The metal fabrication industry is particularly relevant to our intervention since it is an example of knowledge-intensive (and capital-intensive) manufacturing that is concentrated only in

[2] Management quality also matters even among large firms in the developing world (Bloom et al. 2012).

small firms in developing countries. As mentioned, there are thousands of small metal fab fabrication firms in and around Kampala alone, and such firms are typical across the developing world.

4.2.2 Conceptual Framework

Why would firms that stand to increase their profits through technical training not seek such opportunities, or why would workers who could increase their wages not make a profitable investment in the future? Answers to such questions highlight possible frictions that may overwhelm any possible benefits in the absence of our intervention in this context. Anticipation of such frictions may prevent firm owners in our context from demanding training, and thus, training programs may not arise naturally. We summarize these frictions in three categories:

(I) Owners do not have an accurate perception of the value of training.

(II) Owners' selection of workers for training is inconsistent with their beliefs.

(III) Worker's beliefs about the training's value do not align with those of owners.

While these frictions are uniquely tailored to our context, suboptimal investment in human capital is not. Across both the developed and developing world, firms regularly under-appreciate the value of training. We document why these frictions arise and why they could lead to under-investment in training by observing them directly in the context of an exogenously-provided (rather than endogenously-demanded) training program below.

Owners do not have accurate perceptions about the value of training

One hurdle to overcome in establishing a training program could be the existing perceptions about such training programs. In particular, the take-up of a training program will be necessarily hampered by the perception that such programs are ineffective or poorly suited to the needs of the existing marketplace. Moreover, firm owners fail to understand how potential benefits from training may be distributed between the firm and their workers.

Incorrect perceptions about training may lead to under-appreciation of a training program and investment. Education literature shows that lower educated parents may expect lower returns to education, which in turn may affect investment in a child's education (Brown 2006) as beliefs about returns to education can affect investment choices (Attanasio, Boneva, and Rauh 2022). In our setting, firm owners have low levels of formal education and low levels of vocational training, which may affect the effectiveness of our training program. Failure to perceive potential returns to training may be present even among large firms in developing economies, foregoing potential returns to even simple low-cost training (Adhvaryu, Kala, and Nyshadham 2018).

Additionally, owners in our setting may be overconfident in their ability to hire and train their workers, which is strikingly similar to a long history of research on over-confidence among entrepreneurs (Malmendier and Tate 2005). Entrepreneurs may be overly optimistic about their likelihood of success in identifying talented workers or may overstate the abilities of their workforce or their ability to train their workers adequately. This is likely prevalent in our setting as worker roles are flexible, and responsibilities overlap. Thus, overconfident owners may believe that their employees already possess the necessary skills or can learn on the job without formal training. This over-optimism would diminish the demand for training and distort the assignment to training when presented with a free training program.[3]

Owners' selection of workers for training are inconsistent with their beliefs

Even with the correct perceptions about the value of training in terms of worker benefit, firm owners may not select the optimal worker–the most teachable (who would benefit from training the most). Misaligned incentives between owners and workers, behavioral biases, risk aversion, and non-standard production functions could all lead to inconsistencies between beliefs about an individual worker's returns to training and the selection of workers for training, leading owners not to select the most teachable worker. These inconsistencies could manifest themselves in lower average returns to the training program.

Incomplete contracts may characterize this setting. Incomplete contracts naturally lead owners to underinvest in training because of the relationship-specific nature of the investment. This is a classical prediction of the hold-up problem in Hart and Moore (1988) and is elaborated on in Acemoglu and Pischke (1998).[4] Firm owners will seek to assign workers to training from whom

[3]This could also partly explain the prevalence of apprenticeship in Sub-Saharan Africa reported in Filmer and Fox (2014). Also, in addition to the desire to accrue rents from paying less than market wages to low-skilled employees, firm owners may be confident in their ability to train, so they regularly hire low-skilled workers whom they can train and pay less than market wages, at least during apprenticeships.

[4]Moreover, solutions to hold-up problems, such as designing wage contracts to offer different wages for different

they can capture the largest returns. This may involve sacrificing a worker's returns to training for a higher likelihood that the worker remains at the firm.

These firms likely practice relational contracting, implying that owners will rely on other dimensions, such as trust, to select workers for training. This may undermine potential training benefits as this setting is generally characterized by low social trust (Falk et al. 2018). Moreover, social psychology literature shows that non-Western societies emphasize values rooted in loyalty and morals, underscoring relational contracting. Thus, to avoid separation after training, owners may select workers who they believe would stay with the firm based on how much they trust them. Alternatively, firm owners might employ family members or select family members for training because they trust them more.[5]

Alternatively, a profit-maximizing owner might be individually rational, not socially optimal. A socially optimal owner will select the most teachable worker even though only some of the returns from training may be extracted by the owner. An individually rational owner will select the most profitable worker—the worker they perceive to increase the firm's profit the most after training. The most profitable worker may not be the most teachable.

Also, impatience may inhibit the owner from maximizing the returns to training. A present-biased owner may outweigh the loss of production they see from assigning one of their better workers to training and decide to send a lower-quality worker who may not benefit as much. An owner who assigns a worker for training pays an immediate cost in the form of lost productivity from that worker while attending training. This conflicts with the long-run benefit of having a more productive worker after training. A present-biased owner may under-select high-quality workers who may benefit more from training because of the large, immediate loss in productivity. [6]

Lastly, firms with atypical production functions where the complementarity between worker skills is particularly high may rationally select workers for training who will not see the highest individual returns. For instance, if a firm's productivity is limited by the lowest skilled worker,

tasks proposed in Prendergast (1993) and later discussed in Leuven (2005), is difficult. The model in Prendergast (1993) predicts that if firms offer higher wages for difficult tasks, it may induce workers to invest in training to obtain the required skill level. The lack of institutional structures that enforce contracts makes this difficult.

[5]This is consistent with literature (e.g., Bertrand et al. 2008; Ilias 2006) across different settings in the developing world that report high family involvements in micro firms for reasons, such as avoidance of urgency costs.

[6]The analysis presented in the version does not include results on time inconsistency biases of owners. To evaluate the impact of this intertemporal conflict, we will identify whether or not owners who display present-biased preferences over cash payments are more likely to express present-biased preferences in their selection of workers for training.

then the firm owner may maximize their returns by sending the lowest skilled worker for training, regardless of the individual returns. Optimizing under the constraint of such a production function is complicated and requires both correctly perceiving the individual returns to training as well as identifying the complementarities between the skills of workers.

Workers and owners' beliefs about training do not align

Owners and employers may have different objective functions. Individually rational owners seek to maximize profit, while their employees seek to maximize their lifetime earnings (post-training marginal product). The lack of contractual terms that tie an employee's marginal product to training may lead to misaligned incentives for a training program. Just as the owner has the incentive to select workers for training based on their ability to extract the returns from training, workers similarly have the incentive to demand training to maximize private gains from training.

Thus, workers who perceive themselves to be more mobile, more connected to other employers, or more likely to transition to self-employment would have a higher demand for training.[7] Lastly, worker demand for training may be determined by the perceived average improvement from training.

4.3 Research Design

Our study relies on the random assignment of firms to treatment and control groups to study the selection preferences of employees into training according to employer and employee preferences. Appendix 4.8.1 Figure 4.8.1 gives a heuristic process flow of sample construction. Within firms, we assign two workers to training: one according to owner preferences and one according to worker preferences.

Our study relies on two main types of data collection: surveys and practical tests. In addition to detailed information on the firms and workers, the surveys include incentive-compatible elicitation of preferences for training. We also rely on applied practical metal fabrication tests to

[7] Higher demand for training could even show in other forms of training, such as apprenticeships. For example, Frazer (2006) develops a model where even among apprenticeships that are firm-specific, apprentices learn the firm's technology, which they replicate in future self-employment. He confirms the model predictions with the Ghanaian Manufacturing Enterprise Survey and finds that 77 percent of the apprentices express strong preference for self-employment.

measure worker productivity.

4.3.1 Firm Evaluation Sample

We first identify our evaluation sample by relying on a brief census survey of all metal fabrication firms located in Kampala subburbs.[8] We used this screening survey to physically map firm locations and identify firms with the targeted number of employees. We limit eligibility to firms with at least four employees since choosing employees for training is a decision that is relevant to experienced managers who have more than one or two employees. In addition, it would be impractical for owners to lose more than 50% of their workforce during training. We excluded firms where the owner or employees had previously trained with our expert trainer as we did not want this prior interaction to affect take-up. We then approached all the qualifying firms for a marketing exercise and invited the owners for orientation.

We organized marketing to facilitate adequate take-up. During these visits, we talked with both owners and employees since both were to engage in the training either through selection to training or attendance. The marketing visit included an invitation to an orientation event, which was another opportunity for employees and owners to learn more about the training.

Thus, firms are part of the evaluation sample based on these three criteria. First, they passed the number of employees requirement and did not have prior interaction with our expert. Second, they must have attended one of the orientation events. That is, they had concrete information about training to make informed decisions during the incentive selection exercises we discuss below. Lastly, they were located close to one of three training centers.

4.3.2 Sample of Potential Trainees

During the baseline survey, we interviewed the owners and two workers whom we identified to potentially receive training based on the elicitation exercises. Since this elicitation takes place before treatment is assigned, we identify a sample of *potential trainees* during the baseline survey.

We selected the first worker at each firm using workers' collective preferences. Workers from the participating firm cast bids in an auction to be the first potential trainee based on their 'willingness to accept' participation in training. We asked each worker to demand an amount of

[8] We do not list in our pilot area.

between 0 and 120,000 Shillings ($34) they would be willing to accept to attend training. This allowed us to collect information on each worker's value on assignment to training. The worker with the highest value for training will report the lowest willingness to accept for assignment (that is, they require the least subsidy to attend the training), and is thus, selected as the first potential for training. Workers cast bids simultaneously to limit information sharing between workers that would influence bidding behavior, implying we elicit demand for training only if most (not all) of the workers were present at some firms.[9] Since we paid out the demand made by employees whenever they were in the treatment, our demand elicitation was, therefore, incentive-compatible. We call this elicitation "worker demand elicitation".

After worker demand elicitation, we carried out the owner elicitation. The owner's preference determined the second worker of the potential trainees at each firm. The owners revealed their preference for training by ranking their workers from the most preferred worker to the least preferred worker to receive training. Since this ranking determined one of the firm's two potential trainees, it is, therefore, incentive-compatible.

We only revealed the winner of worker demand elicitation simultaneously with the owner's election after the owner elicitation. We did this to ensure that the owners' true ranking was not influenced by their knowledge of the winner of the worker demand elicitation exercise at each firm. In case of a tie between owner selection and demand elicitation, we pick the second preferred worker by the owner and the winner of the demand elicitation. The ties happened about 20 percent of the time during our data collection.

4.3.3 Training Treatment

We randomly assigned firms in our evaluation sample where the baseline surveys (owner and employee) and practical tests had been completed into treatment and control groups. The main training intervention for this study focuses on technical training for employees of small firms. We invested substantially in developing this technical training program for this study with the intention that it could be implemented more widely once the study is complete. We identified a metal fabrica-

[9] There were cases where not all firm employees were present during this elicitation exercise. In such cases, we only carried out elicitation if a sufficient percent of the workers were present at the firm during the survey date, otherwise, we would reschedule the survey for the following day. Specifically, we required at least three out of four or five-employee firms, 4 out of 6-8 employee firms, and at least five employees for larger firms to be present before eliciting demand.

tion expert who, before designing the training, conducted semi-structured interviews with firms to determine the gaps in their technical knowledge. After the training was initially piloted, we hired a curriculum expert to work with our implementing expert to develop a detailed and potentially scalable lesson-by-lesson curriculum.[10]

We designed the training to engage and be accessible to employees of small manufacturing firms. This population has limited formal education and has low levels of literacy, as Table C2 shows. This is in contrast to publicly available metal fabrication training in Kampala, which often expects trainees to have had a high school education. Thus, the training focuses on demonstrations and hands-on practice and does not expect trainees to engage with detailed written materials.

The design of the training itself also accounts for a target population that is very low-income and has a high opportunity cost of time since they are already working. Thus, the training took place in temporary training facilities that are close to the employees' workplaces to facilitate regular attendance. Thus, to ensure workers are not away from their firms for more than a third of their working days per week, we conducted no charge training sessions just two days a week for six to eight hours a day, for a total of 23 days (totaling 168 hours) for each training cohort. Additionally, all trainees received meals and transport reimbursements on training days.

4.4 Data and Outcomes

We collected data at three points in time. At baseline, we surveyed owners and workers. These surveys include pre-specified incentive-compatible preference elicitation and the selection of workers in the potential trainee sample in Section 4.3.2 described above. Also, at baseline, we implemented practical (i.e. applied metal fabrication) tests for the potential trainee sample. Then, immediately after completing training, we conducted a follow-up practical test.

4.4.1 Baseline Survey–Owners

Once we identified the evaluation sample, we conducted a baseline survey of the owners of all participating firms. We collected data on the owner, the firm, and the firm's employees. The information

[10]The intention is that this curriculum could be implemented by skilled metal fabricators with limited prior teaching experience. Scalable curriculums are increasingly common in development, especially for business skills training. To our knowledge, this is the first technical skills training that is designed to be scaled. Moreover, this training curriculum is at par with the expectations of Uganda's Level One Directorate of Industrial Training (DIT) for metal fabrication.

on the firm includes data on profitability, assets, access to credit, some measures of productivity, and personnel practices. The survey also included a worker roster that collected information on a worker's history with the firm as well as their reliability and productivity (as perceived by the owner). This will be our primary source of data and includes all of the metal fabrication workers at the firm. The information from owners includes incentive-compatible elicitation of the owner's time preference and a measure of risk attitudes.

In addition to the incentive-compatible selection for training, the essential information from the baseline survey is the incentive-compatible elicitation to capture the owner's beliefs about their worker's quality and gains of training for each of their workers. In these elicitations, owners provided a guess of scores each of the workers would obtain in an objective applied metal fabrication test–practical tests as described in Section 4.4.3 below. Specifically, we ask about beliefs about their workers' scores on this practical test that would be scored on a 0-100 scale.[11] Specifically, we ask owners about their beliefs about each of their workers' competency: (I) at baseline, (II)) at follow-up, assuming they are not trained, and (III) assuming that they are trained. The difference between (III) and (II) is our pre-specified measure of <u>teachability</u>. We pay owners whenever their beliefs are within the 10 percent range of the correct score.

In addition, we ask the owners about the perceived short-run opportunity cost of sending workers to training and the perceived post-training profitability of training each worker. That is, each owner ranks each of their workers in terms of perceived profitability gain. These data from elicitations and perceived profitability form the central part of the analysis, together with practical skills test scores.

4.4.2 Baseline Survey–Workers

As part of the baseline survey, we interview the two potential trainees from each firm once they are identified. The worker roster in the owner survey collects basic information on each worker, such as their tasks and their work history at the firm. The worker survey, however, collects a more detailed work history as well as complete wage information, particularly from any work done at other firms. The survey also collects more in-depth data on worker characteristics, mainly Raven's matrices, to

[11] When doing this, we explain to the owners that a score of fifty would correspond to an average metal fabricator in Kampala, a score of zero would correspond to a person without any metal skills, and a score of 100 would correspond to a metal fabrication expert.

capture cognitive ability. In addition, we will conduct incentive-compatible elicitations of risk and time preferences.

We asked workers about their beliefs about the impact of training. The question asking owners about their beliefs on the perceived impact of our training on their competency was not incentivized to prevent hedging behavior. We incentivized, however, the question about their belief about the average score of the trained worker. Lastly, we asked workers about how they think gains from training would be shared between the firm and themselves (i.e., if their wage will be increased after training).

4.4.3 Practical Skills Tests

After we had determined the two potential trainees from each firm, we invited all potential trainees to participate in a baseline practical skills test. The skills test involved the fabrication of a standard metalworking product in such a way that evaluators could test several relevant skills (e.g., measuring, cutting, welding, grinding). The evaluators assessed each worker based on quality, speed, safety, and the efficient use of materials. Scores range from 0 to 100 percent, with the highest score corresponding to the output of a master metal fabricator with decades of experience. The tests were scored by accredited assessors (different from the trainer) who teach at one of the prestigious and accredited vocational institutes in Kampala.

Soon after concluding the training, potential trainees from the treatment and control groups repeated the baseline practical skills test. With repeated observations, we can calculate the improvement of each worker and evaluate the effect of the training on technical skills. We can then compare true improvement in skills to the perceived improvement collected in our incentive-compatible elicitations.

4.5 Preliminary Results

Our experiment is still ongoing at the endline practical test stage. The results I present in this section are preliminary and correlational. We provide regression equations for the results in each respective section.[12]

[12] We shall revise this to include the empirical strategy section when updating the analysis with endline data

4.5.1 Descriptive Statisitics

Table C1 reports firm-level summary statistics, while Table C2 reports summary statistics of the workers in the potential trainee sample. Our sample comprises dynamic firms and active entrepreneurs in terms of employee and product turnover who have been in business for an average of ten years. Additionally, an average of one employee quit, while an average of two new employees were hired in the firm's previous year. Despite the artisanal business nature of the firms in our setting, close to 50 percent of the entrepreneurs reported that they introduced new products in the last year. The average profit was approximately 1,551,100 Ugandan shillings ($430).

The average number of employees is five, which is comparable to the 5.88 sample average of welding firms in the same setting reported in Bassi et al. (2023). Almost all the entrepreneurs in our sample are committed, not survival entrepreneurs, as only six percent mentioned they would leave metal fabrication if offered a job with a wage equivalent to their current profits. Also, these entrepreneurs are optimistic about the future. More than 80 percent expect their business to survive for at least five years, and the average number of employees is expected to more than double in the next five years. Additionally, these entrepreneurs expect close to 70 percent of current businesses to survive for at least five years, portraying a general positive outlook on the economy.

These entrepreneurs have low levels of formal education. The average years of formal education is 10.4, equivalent to completing Year Three of O-level (Ordinary Level) education in secondary school. Moreover, less than 20 percent have completed formal metal fabrication training. Thus, most of these entrepreneurs likely learned their skills on the job, as most have been entrepreneurs for more than ten years. It is possible these entrepreneurs may not correctly perceive returns to training.

Credit constraints may be a factor affecting the provision of training in our setting, as less than five percent have applied for loans to train workers. When we asked why they had not applied for training, most of these entrepreneurs believed it was unnecessary, while others mentioned they did not know how or where to apply for loans. It is likely that entrepreneurs in our sample rely on their own ability to train their staff. It is thus unsurprising that more than 80 percent of employees mentioned that they had been apprentices at some point at the current firm, as Table C2 shows.

Turning to Table C2, the average age of workers is 24 years old, and all the workers employed in this subsector are male. Additionally, it appears most of the workers in our sample started at their current firm, as the ratio of years in metal fabrication and years at their current firm is 1.01. As with the owners, their education level is low. The average education is barely one year of post-primary education, which is two years less than the average of owners reported in Table C1. Moreover, only 5 percent of the employees have completed vocational training, indicating a need for training programs Yet, Alfonsi et al. (2020) show vocational training in this setting is more beneficial than on-the-job training.

The average hourly wage is 4,140 shs ($1.14), and our sample comprises full-time employees, defined as working at least four days per week at the firm, as per the sample restriction. Although not reported, most of the employees are paid based on completed orders. Additionally, a quarter of the employees have worked for more than one firm, and close to half discuss job opportunities with someone from a different firm. Lastly, more than 80 percent stated that they had ever been apprentices at the firm, implying that most workers use apprenticeships to obtain employment in the sector and that most employees are engaged in some informal training. The average in this sector is way larger than the proportion of apprentices in other sectors Filmer and Fox (2014).[13]

4.5.2 Beliefs about Worker Quality

Figure B1 plots the owner baseline perception of quality against the workers' practical test scores, which is our objective measure of quality. Additionally, the plot includes three lines of fit: the 45-degree line, the best-fit for owner-selected workers, and the best-fit for other workers. We can note several findings from this figure. First, owners are sophisticated in the sense that they can distinguish a high-quality worker from a low-quality worker among their metal fabrication staff. The correlation between perceived quality and objective score is high (close to 70 percent).

Nevertheless, owners seem to overestimate the quality of their workers on average. From the Y-axis, there are several firms where the perceived quality is above 80 percent or even equal to 100, but none of the workers scored above 80 percent in the baseline practical test. Interestingly, owners seem to perceive their selected workers to be of higher quality than the other workers. Yet, the

[13]Filmer and Fox (2014) Figure 3.18 reports that an average of 20 percent of SSA youth have ever been an apprentice, where the proportion is close to eight percent in Uganda. They report statistics of apprenticeships across all sectors in six countries using a standardized survey. This implies Uganda's metal fabrication subsector may be an outlier, as apprenticeships are a more prevalent way of finding employment and acquiring metal fabrication than other sectors.

correlation of their perceptions with practical tests is lower than that of worker-selected workers. That is, the slope of the line of best-fit for owner-selected workers is lower than the slope of the line of best-fit for all other workers.

4.5.3 Do Owners Select the Most Teachable Worker for Training?

Figure B2 plots the perceived benefit to workers from training each worker on the Y-axis. Each spike represents a firm, and each dot represents the perceived teachability of each worker. The longer the spike, the larger the dispersion in perceived teachability within a firm. From the graph, perceived teachability is positive on average. That is, most owners believe that their workers' quality will improve with training. Despite their beliefs, a lot of owners do not select the worker that would improve most from training, implying that owners are not socially optimal.

4.5.4 Do Owners Select the Most Profitable Worker for Training?

Figure B3 plots owner selection for training on the left panel and most teachable on the right panel by rank of perceived profitability post-training on the right panel. Suppose i represents a worker, and k represents the ranking of perceived profitability. Each bar on the left panel can be read as "What proportion of the owners that selected worker i ranked k in terms of perception of how much i would each increase a firm's profit after training?" Using our teachability measure, we created a dummy that is equal to one when a worker is the most teachable (right panel) at a firm and zero otherwise. Thus, we interpret each bar on the right panel as "What is the proportion of firms where the most teachable worker ranked k in terms of perception of how much profit after training?"

From both panels, we see that owners are sophisticated in that only a small proportion cannot rank their workers in terms of profitability. From the left panel, we observe that more than 70 percent of owners selected workers that ranked first or second in terms of profitability for training, and close to 55 percent selected the worker they perceived to bring in the most profit to the firm for training.

In the right panel, there seems to be no correlation between perceived profitability and perceived teachability. That is, most owners do not perceive the most profitable as the most teachable, as only about 10 percent perceive the most teachable worker as the most profitable. Additionally, a significant number of firms perceive the least profitable workers as the most teachable, with close to 40 percent ranking the most teachable fourth or greater.

Taken together, these results show that owners are individually rational in selecting workers to maximize profits. However, they do not perceive this increase in profitability to come from workers who would improve most from training. This implies that owners consider other dimensions in their decision-making, such as considering which workers might separate after training. We report these results in Section 4.5.5 below.

4.5.5 Employee Firm Ties and Selection for Training

To study what variables predict owner selection and perceived profitability, we use the following regression:

$$\text{(A1)} \qquad Y_{ij} = \beta_0 + \phi_1 Family_{ij} + \phi_2 Trust_{ij} + \phi_3 Separation_{ij} + \phi_4 Tenure_{ij} + \delta_j + \varepsilon_{ij},$$

where Y_{ij} is the outcome of interest, such as owner j's selection ranking of worker i's for training. $Family_{ij}$ is an indicator that is equal to one if the worker is related to the owner, $Separation_{ij}$ is the perceived likelihood that the worker will leave the firm in the one year, and $Tenure_{ij}$ is how long the firm has employed a worker, and δ_j are firm fixed effects.

Table C3 presents several versions of equation (A1) where the selection ranking of workers for training is the outcome in (1) and (4), perceived profitability in (2) and (5), and teachability in column (3). Additionally, columns (4) and (5) control for the percentile of a worker's perceived teachability. All regressions control for fixed effects, and standard errors are clustered at the firm level. These regressions can be interpreted as correlational rather than causal, as the experiment is ongoing.

The findings reveal that the owner consistently selects a worker who is a relative for training. This coefficient is positive and significant at the one percent level in (1) and does not change when we control for perceived teachability in (4). Additionally, a family member is perceived to be more profitable. The coefficient is also positive and significant in both columns (2) and (5). A coefficient of 0.101 from (1) indicates a family member is ranked 10.1 percentiles higher compared to a non-family member. Also, an owner perceives a family member to increase a firm's profitability by 10.3 percentiles more than non-family members when selected for training. Lastly, being a relative is not associated with perceived improvement from training.

Trust, measured by worker perceived reliability, is also positively and highly significant when the outcome is owner selection for training or perceived profitability. From column (1), a one-unit increase in trust, measured on a Likert scale, is associated with an increase of about 8.1 percentiles in the worker's ranking for selection for training, while the same change in trust is associated with a 10.4 percentile increase in profitability in column (2). On the contrary, trust and perceived teachability are negatively correlated. A one-unit increase in trust is associated with a decrease of seven percentiles in teachability.

Conversely, the likelihood of leaving is negatively associated with training selection or perceived profit and is not significant when the outcome is teachability. From (1), a one-unit increase in the perceived likelihood of worker separation in one year is associated with a decrease of approximately one percentile in their training selection ranking or perceived profitability.

Tenure shows mixed results. The coefficient is negative and significant in (3) when teachability is an outcome, barely significant when the outcome is selection for training, and not significant when the outcome is perceived profitability. From (3), one more year of tenure at the firm is associated with a 1.4 percentile decrease in teachability and a 0.6 percentile less likely selection for training. This result implies that owners may have confidence in their ability to train. They may perceive workers who have been at the firm for a long time to have already gained skills such that additional training would not be beneficial.

Lastly, we add teachability as a control in models (4) and (5), which shows a negative and significant impact in model (5) when perceived profitability is an outcome but not significant in model (4) when selection for training is an outcome. The results from (5) indicate that a one-percentile increase in teachability is associated with a 0.1 percentile decrease in perceived profitability. This counterintuitive result suggests that workers who are perceived as more teachable are also perceived to be less profitable, implying factors that proxy relational contracting, such as family, trust, and perceived separation, matter for owner selection for training and this perceived profitability.

4.5.6 Worker Demand

To study what variables that predict worker demand for training, we use the following regression:

(A2)

$$D_{ij} = \beta_0 + \phi_1 Quality_{ij} + \phi_2 WGain_{ij} + \phi_3 OGain_{ij} + \phi_4 WageChange_{ij} + \phi_5 ExpectedTenure_{ij} + \delta_j + \varepsilon_{ij},$$

where D_{ij} is worker i's at firm j demand for training. $Quality_{ij}$ is the perceived quality of the trained staff at the endline, $WGain_{ij}$ the worker's own perceived gain in quality from training, $OGain_{ij}$ is the owner's perceived gain in quality from worker training, $WageChange_{ij}$ is the perceived wage change after training, $ExpectedTenure_{ij}$ is he expected tenure at the current firm, and δ_j are firm fixed effects.

The perceived average score of trainees reflects the incentivized average score of trained staff on the endline practical test. Own perceived benefit is the measure of how much a worker believes he would gain in skill if selected for training, which may be different from the owner's perception of the benefit.

The outcome variable is the percentile of demand for training. All regressions control for fixed effects. The difference between columns one and two is from the variables we control for. We do not have data on perceived wage change after training for all the workers, and some workers did not know how long they expected to stay at the firm. Therefore, the number of observations in column two is reduced compared to column one. The regressions from columns one and two show that the owner's perceived teachability training is the strongest predictor of workers' demand for training. Additionally, the perceived gain for training is associated with an increase in demand for training. However, the coefficient of own perceived benefit is qualitatively different between columns one and two.

4.5.7 Worker Demand Vs Owner Selection

Figure B4 is a scatter plot of the owner's percentile ranking for training on the Y-axis and the percentile of worker demand for training on the X-axis. As expected, we observe bunching around one on both axes. The worker selected for training is consistently ranked number one, hence the bunching around one on the Y-axis. Similarly, the worker with the lowest demand for training corresponds to the highest rank/percentile, leading to bunching around one on the X-axis. The graph indicates weak alignment between worker demand for training and the owner's selection for training, implying that the workers selected by the owner are not necessarily those who demand training the most.

4.6 Conclusion

The classic Becker (1962) model of human capital predicts that firms will underinvest in general skills training because, in competitive labor markets, firms will pay employees less than their marginal product to recoup training costs, prompting employees to leave for other firms. This underprovision is especially pertinent in developing countries, where low productivity constrains growth. Government subsidies for training are common in richer countries to address this issue, but in developing countries, information frictions may lead firms to select suboptimal employees for training even if it is subsidized. This study examines the suboptimal provision of employee training through an experiment in Uganda's metal fabrication sector, offering free training and analyzing whether owners choose the socially optimal worker who would maximize sectoral productivity or the one who maximizes firm profits but would not improve most from training.

We conducted data collection with metal fabrication firm owners, including incentive-compatible selection for training. We also elicited owners' perceived quality of their workers at baseline and endline, both with and without training, as well as perceived profitability from training each worker and data on worker ties to the firm, such as kinship, trust, and perceived likelihood of separation. From the workers' side, we gathered incentive-compatible demand for training by asking them to specify a willingness-to-accept (WTA) amount to attend training, with winners paid their WTA if selected. Workers' quality was objectively assessed through practical tests scored by vocational institute assessors. We randomized small metal fabrication firms (4-14 employees) into treatment and control groups, offering a training program to the treatment group, using a curriculum designed with input from metal fabrication experts and vocational lecturers.

Our preliminary analysis reveals that, on average, firm owners believe that our training program positively impacts their workers' quality. However, they tend to select workers based on perceived post-training profitability rather than those who would maximize productivity gains for the metal fabrication sector. Owners prefer workers they trust and those with strong ties to the firm, such as family members, ranking them significantly higher for training. Despite free training eliminating credit constraints and mitigating the anticipated risk of separation since owners can afford to pay workers their post-training marginal product, owners remain individually rational, choosing workers with the largest gap between post-training marginal product and wage, often relatives or highly trusted workers.

Our findings suggest that interventions to improve the effectiveness of training programs should consider frictions, such as separation, and aim to align owner-worker incentives better. Lastly, while we find that owners have positive perceptions of our training program and training in general, the analysis at this stage does not test whether they are underestimating or overestimating the effectiveness of training. We shall test this once our endline is complete. We find, however, that the owners are overestimating their worker quality on average.

References

Acemoglu, D. 1997. "Training and Innovation in an Imperfect Labour Market." The Review of Economic Studies 64:445–464.

Acemoglu, D., and J.S. Pischke. 1998. "Why Do Firms Train? Theory and Evidence." The Quarterly Journal of Economics 113:79–119.

Adhvaryu, A., N. Kala, and A. Nyshadham. 2018. "The Skills to Pay the Bills: Returns to On-the-job Soft Skills Training." Working Paper No. 24313, National Bureau of Economic Research, February.

Alfonsi, L., O. Bandiera, V. Bassi, R. Burgess, I. Rasul, M. Sulaiman, and A. Vitali. 2020. "Tackling Youth Unemployment: Evidence From a Labor Market Experiment in Uganda." Econometrica 88:2369–2414.

Attanasio, O., T. Boneva, and C. Rauh. 2022. "Parental Beliefs about Returns to Different Types of Investments in School Children." Journal of Human Resources 57:1789–1825.

Bassi, V., J.H. Lee, A. Peter, T. Porzio, R. Sen, and E. Tugume. 2023. "Self-Employment Within the Firm." Working Paper No. 31740, National Bureau of Economic Research, September.

Becker, G.S. 1962. "Investment in Human Capital: A Theoretical Analysis." Journal of Political Economy 70:9–49.

Bertrand, M., S. Johnson, K. Samphantharak, and A. Schoar. 2008. "Mixing family with business: A study of Thai business groups and the families behind them." Journal of Financial Economics 88:466–498, Darden - JFE Conference Volume: Capital Raising in Emerging Economies.

Bloom, N., B. Eifert, A. Mahajan, D. McKenzie, and J. Roberts. 2012. " Does Management Matter? Evidence from India *." The Quarterly Journal of Economics 128:1–51.

Brown, G., M. Hardy, I. Mbiti, J. McCasland, and I. Salcher. 2024. "Can Financial Incentives to Firms Improve Apprenticeship Training? Experimental Evidence from Ghana." American Economic Review: Insights 6:120–36.

Brown, P. 2006. "Parental Education and Investment in Children's Human Capital in Rural China." Economic Development and Cultural Change 54:759–789.

Card, D., P. Ibarrarán, F. Regalia, D. Rosas-Shady, and Y. Soares. 2011. "The Labor Market Impacts of Youth Training in the Dominican Republic." Journal of Labor Economics 29:267–300.

Cefala, L., P. Naso, M. Ndayikeza, and N. Swanson. 2023. "Under-training by Employers in Spot Labor Markets: Evidence from Burundi." Unpublished, Mimeograph.

Cho, Y., D. Kalomba, A.M. Mobarak, and V. Orozco. 2013. Gender Differences in the Effects of Vocational Training: Constraints on Women and Drop-Out Behavior. The World Bank.

Falk, A., A. Becker, T. Dohmen, B. Enke, D. Huffman, and U. Sunde. 2018. "Global Evidence on Economic Preferences*." The Quarterly Journal of Economics 133:1645–1692.

Filmer, D., and L. Fox. 2014. Youth Employment in Sub-Saharan Africa. Africa Development Series, Washington, DC: World Bank.

Frazer, G. 2006. “Learning the master’s trade: Apprenticeship and human capital in Ghana.” Journal of Development Economics 81:259–298.

GoU. 2015a. “Second National Development Plan (NDPII) 2015/16 – 2019/20.” Working paper, Government of Uganda, National Planning Authority.

—. 2018. “Trade and Industry Sector Statistical Abstract.” Working paper, Ministry of Trade, Industry and Cooperatives (MTIC), Kampala, Uganda.

—. 2015b. “Uganda Micro, Small and Medium Enterprise (MSME) Policy: Sustainable MSMEs for Wealth Creation and Socio-Economic Transformation.” Working paper, Ministry of Trade, Industry and Cooperatives (MTIC), Kampala, Uganda, June.

Hart, O., and J. Moore. 1988. “Incomplete Contracts and Renegotiation.” Econometrica 56:755–785.

Hirshleifer, S., D. McKenzie, R. Almeida, and C. Ridao-Cano. 2016. “The Impact of Vocational Training for the Unemployed: Experimental Evidence from Turkey.” The Economic Journal 126:2115–2146.

Hsieh, C.T., and P.J. Klenow. 2009. “Misallocation and Manufacturing TFP in China and India*.” The Quarterly Journal of Economics 124:1403–1448.

Ilias, N. 2006. “Families and firms: Agency costs and labor market imperfections in Sialkot’s surgical industry.” Journal of Development Economics 80:329–349.

Leuven, E. 2005. “The Economics of Private Sector Training: A Survey of the Literature.” Journal of Economic Surveys 19:91–111.

Malmendier, U., and G. Tate. 2005. “CEO Overconfidence and Corporate Investment.” The Journal of Finance 60:2661–2700.

McKenzie, D. 2017. “How Effective Are Active Labor Market Policies in Developing Countries? A Critical Review of Recent Evidence.” The World Bank Research Observer 32:127–154.

McKenzie, D., and C. Woodruff. 2016. “Business Practices in Small Firms in Developing Countries.” Management Science 63:2967–2981.

Olafsen, E., and P.A. Cook. 2016. Growth Entrepreneurship in Developing Countries: A Preliminary Literature Review. Washington, DC: The World Bank Group, License: Creative Commons Attribution CC BY 3.0.

Prendergast, C. 1993. “The Role of Promotion in Inducing Specific Human Capital Acquisition.” The Quarterly Journal of Economics 108:523–534.

UBOS. 2014. “Statistical Abstract. Uganda Bureau of Statistics.”

Van Biesebroeck, J. 2005. “Firm size matters: Growth and productivity growth in African manufacturing.” Economic Development and cultural change 53:545–583.

4.7 List of Figures and Tables

4.7.1 Figures

Figure B1: Owners Overestimate Worker Quality.

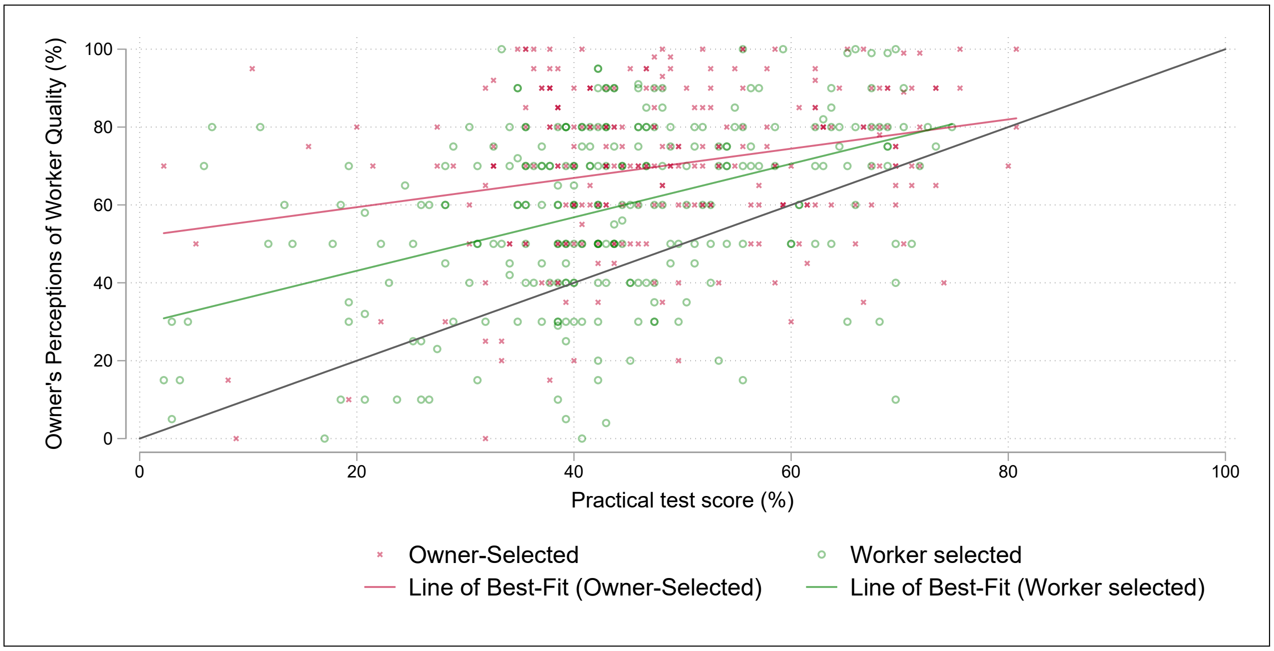

The X-axis (practical test score) is the objective measure of baseline quality, while the Y-axis is the perceived measure of quality. As mentioned in the data section, the perceived baseline was incentivized. We surveyed each owner about what they thought each of their workers would obtain if they were invited to take part in the test, and we gave an owner 10,000 shillings whenever their guess was within the correct range.

Figure B2: Teachability: Owner Selection vs Other Workers

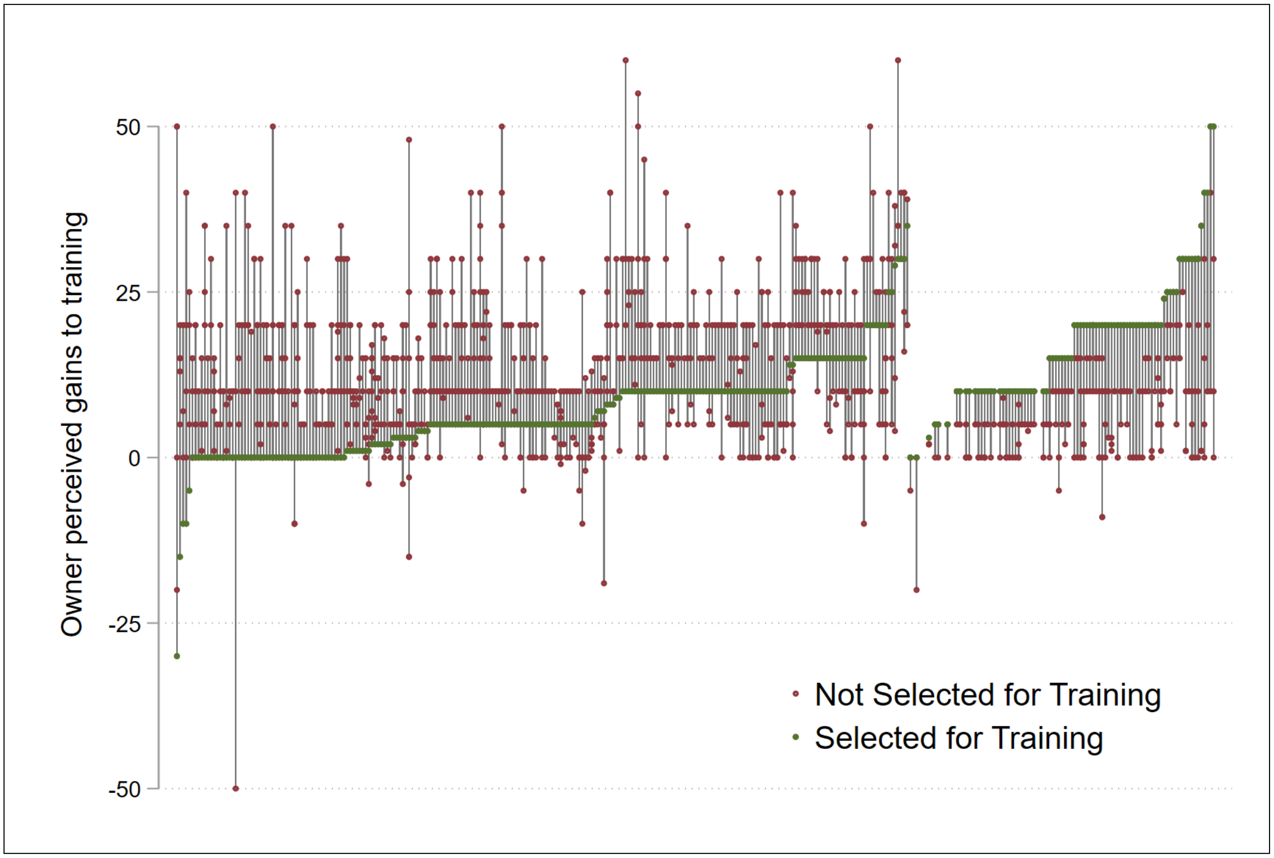

Notes. Each line/spike represents a firm. The firms are sorted first on the perceived gains of the selected worker and second on the percentile rank of workers selected for training relative to other workers within the firm. We exclude cases when the within-firm standard deviation in the perceived teachability is zero.

Figure B3: Perceived Profitability: owner Selection vs Most Teachable

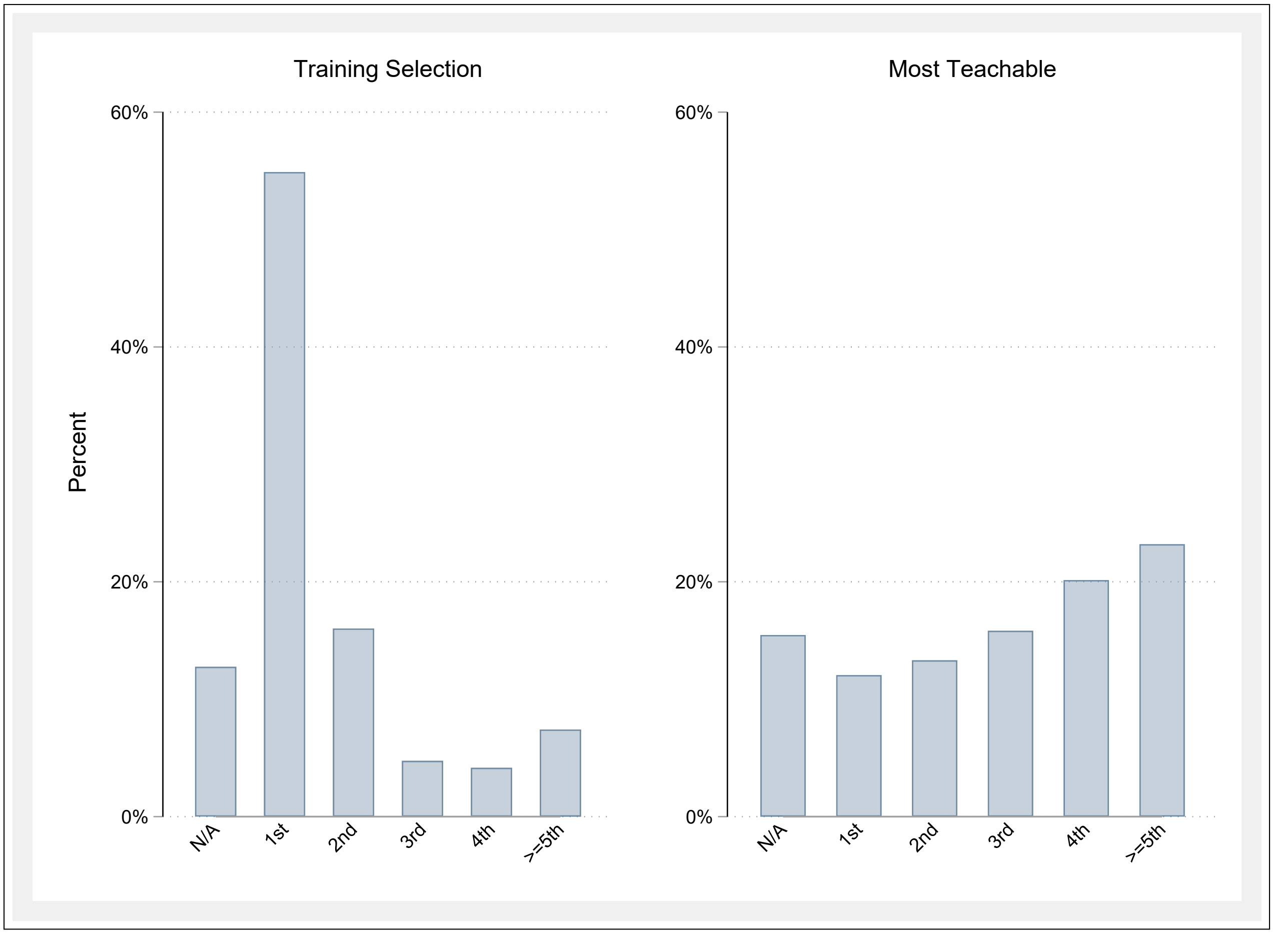

The X-axis is the rank in terms of perceived profitability. We surveyed each owner to determine which of their workers would lead to a profit increase after they had been trained. After this, each owner ranked the workers that would lead to profit from the most profitable to the least profitable. "Ranked None" refers to owners who believe that training does not increase profits or who cannot differentiate the profitability of their workers.

Figure B4: Perceived profitability: Owner selection vs worker demand most teachable

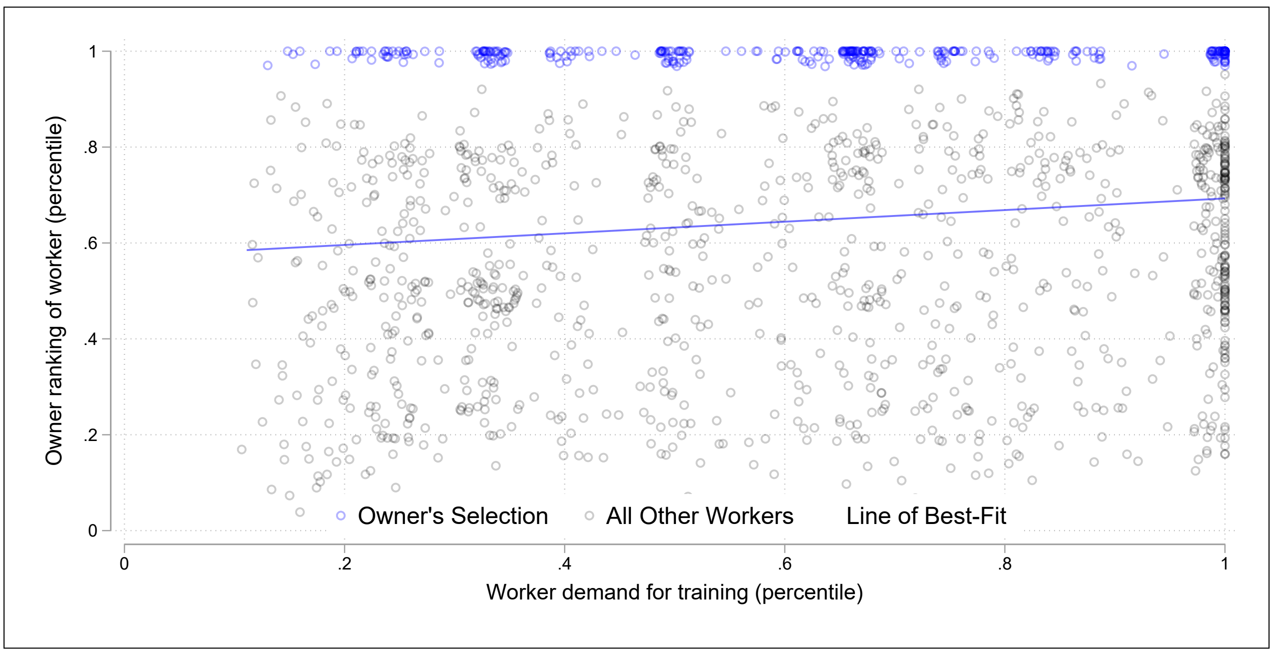

Bunching on one is expected as the selected workers will rank on top of the distribution.

4.7.2 Tables

Table C1: Owner/Firm-level Descriptive Statistics

	N	Mean	SD	Min	Max
Age	336	36.50	8.85	19	73
Years in metal fabrication	337	14.25	7.86	0	49
Completed formal metal fabrication training	275	0.19	0.40	0	1
Education (years)	336	10.35	4.18	0	22
Would shut down business a job offer(wage=profits)	337	0.06	0.23	0	1
Firm age	330	10.35	7.95	0	51
Number of employees	336	5.35	1.82	4	14
New employees hired last 12 months	336	2.26	3.40	0	30
Number of Employees who quit last 12 months	337	1.02	1.72	0	12
Introduced new products last 12 months	337	0.45	0.50	0	1
Revenue in the last 30 days (10,000)	337	803.35	1884.09	0	30,000
Profits in the last 30 days (10,000)	337	155.11	224.44	0	1,500
Employees specialize on tasks	337	0.29	0.46	0	1
Very Likely still in business in 5 years	337	0.84	0.36	0	1
Expected number of employees in 5 years	333	11.41	8.26	2	70
Expected number of similar firms (out of 10) in 5 years	333	6.88	2.61	1	10
Ever received a loan	337	0.25	0.43	0	1
Ever applied for a loan to train employees	337	0.03	0.16	0	1
Ever received for a loan to train employees	337	0.01	0.11	0	1
Change wages after working	337	0.92	0.28	0	1

Notes: Data are from the baseline survey. Revenue and profit are measured in Ugandan Shillings (shs). To convert shs to dollars, divide the shs amount by 3,600. Unless specified, the variable is binary whenever the min is zero and the max is one, such as "Completed formal metal fabrication training". Ever received loan refers to any other type of loan irrespective of training loan. We interviewed businesses if they have ever received any loan from any commercial bank, microfinance institution, savings cooperative, or other certified institutions separately from if they have ever applied or received a loan to train their workers specifically.

Table C2: Worker-level Descriptive Statistics

	N	Mean	SD	Min	Max
Age	674	24.07	5.19	18	53
Female	674	0.01	0.10	0	1
Years in metal fabrication	673	4.53	3.76	0	23
Years at firm	573	4.49	5.04	0	50
Education	673	8.34	3.14	0	22
Completed vocational training	674	0.04	0.20	0	1
Ever under apprentice at firm	674	0.82	0.39	0	1
Hourly wage	660	4139.50	3300.13	0	12,373
Days per week	674	5.79	1.26	0	7
Days last month	673	22.58	7.36	0	30
Worked for other firm last month	604	0.24	0.43	0	1
Discusses opportunities with different firm workers	674	0.49	0.50	0	1

Notes: Data are from the baseline survey. Revenue and profit are measured in Ugandan Shillings. An average education of 8.34 years is equivalent to completing one and a half years of secondary education. Divide by 3,600 to convert the hourly wage to dollars. A few of the workers did not remember when they started at the firm. Ever under apprentice at the firm includes both employees who were currently at the firm under apprenticeship or started and completed an apprenticeship at the firm. Opportunities refer to job opportunities. We asked workers if they share/discuss job opportunities with workers from different firms.

Table C3: Factors Affecting Selection for Training

	(1) Owner ranking	(2) Profitability ranking	(3) Teachability Pctile	(4) Owner ranking	(5) Profitability ranking
Family	0.101***	0.103***	0.007	0.100***	0.104***
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Trust	0.083***	0.104***	-0.070***	0.081***	0.097***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Likelihood of leaving	-0.011***	-0.010**	-0.002	-0.011**	-0.010**
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Tenure	-0.006*	0.004	-0.014***	-0.006*	0.003
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Teachability percentile				-0.033	-0.093***
				(0.03)	(0.03)
R-squared	0.085	0.148	0.096	0.086	0.155
N	1,782	1,722	1,776	1,776	1,714

The outcomes in columns are 0-1 percentile range of the indicated variable. Family is an indicator of whether a worker is related to the owner. Trust and likelihood of leaving are captured using a 0-10 Likert scale. Trust measures the extent to which the owner can depend (all the time) on the worker to do assigned work reliably, while the likelihood of leaving measures the perceived separation of a worker in on year.
All regressions control for control for firm FEs.
Standard errors in parentheses
* p<0.10, ** p<0.05, *** p<0.01

Table C4: Worker Demand for training

	(1)	(2)
Perceived average score of the trainees	-0.161	-0.166
	(0.17)	(0.20)
Own perceived benefit from training	0.051**	0.496*
	(0.02)	(0.28)
Owner perceived benefit from training	0.874***	0.736***
	(0.19)	(0.27)
Perceived wage change		0.003
		(0.01)
Expected tenure		-0.003
		(0.00)
R-squared	0.053	0.063
N	636	495

The outcome is a percentile of worker demand for training, which we measured using the incentive-compatible amount of money a worker was willing to accept to attend training. Expected Tenure is measured in years, while all other controls are measured on a 0-1 scale.
All regressions control for control for firm FEs.
Standard errors in parentheses
* p<0.10, ** p<0.05, *** p<0.01

4.8 Appendix

4.8.1 Time Line and Sample Construction

Our research and data collection process began with an in-person census of mental fabrication firms in the Kampala metropolitan areas coupled with marketing. Our two-person team (enumerator and marketer) visited all firms in our study area, identified those that meet the criteria, and introduced the program. Our marketers were students from nearby vocational schools with the ability to clearly explain the nature of the training, including potential skills and benefits, with the purpose of recruiting businesses in the study area into our training program. We gave business owners an opportunity to opt in by applying.

One month later, business owners who applied and their workers were invited for a 3-4 hour orientation where they interacted with our implementing partner and accredited trainers. The trainers provided a high-level introduction to the curriculum and other training details, such as facilitation on training days. Both listing and orientation were crucial in identifying our evaluation sample by identifying firms that met the number of employees requirement or were located at the training centers.

Figure 4.8.1 highlights the workflow of the sample construction. Every firm that attended the orientation had more information about our training program, making the owner likely to buy into our training program. We also considered the firms located close to the training centers even though they did not attend orientation, as these would incur a small cost (in terms of travel time). As mentioned, our training program covered the financial cost eliminating credit constraints that would affect employee training. We randomized at the firm level after conducting the post-baseline survey. As the figure shows, two employees are selected at each firm through incentive-compatible demand elicitation with the workers or owner elicitation for worker selection. We identify spillover workers as any other workers at the treatment firm whom the owner did not select for training or did not have the highest value for training.

Figure 4.8.1 illustrates our sample selection procedure:

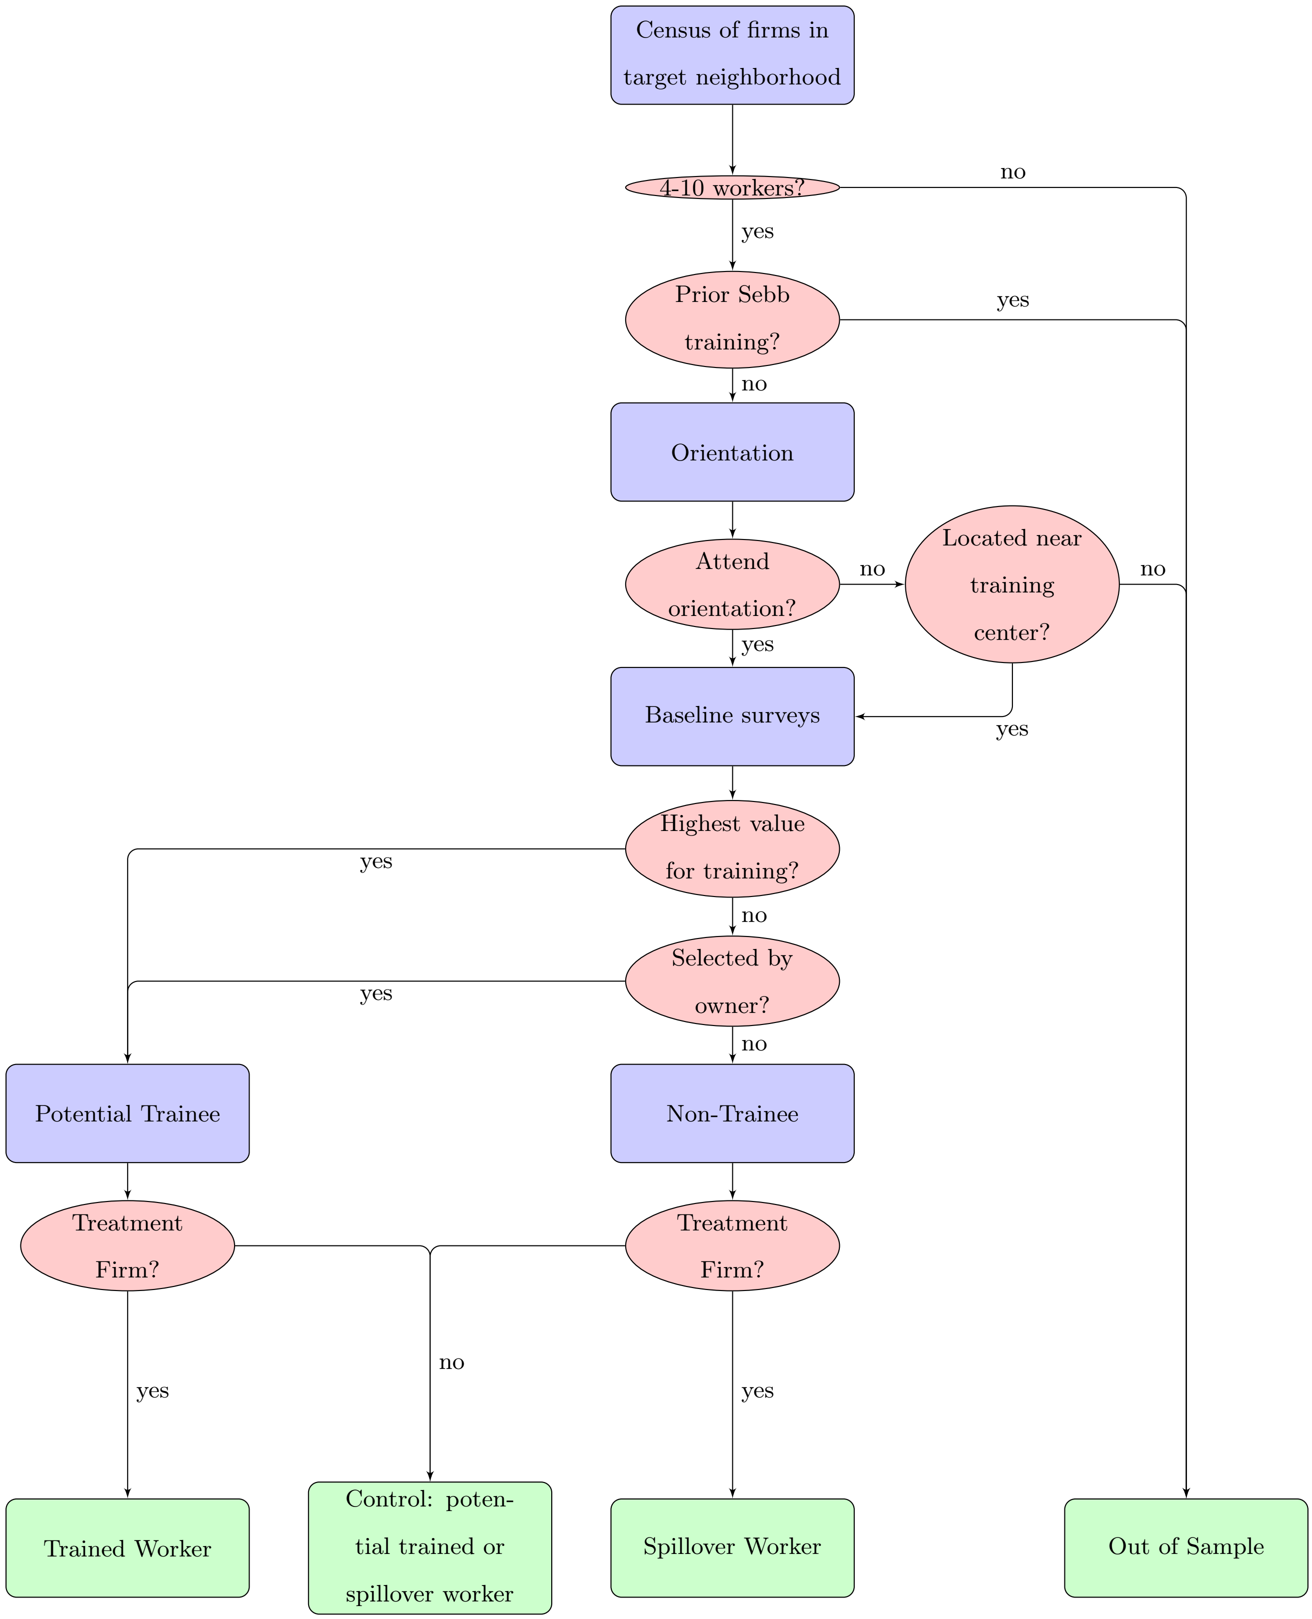
Census of firms in target neighborhood
4-10 workers?
no
yes
Prior Sebb training?
yes
no
Orientation
Attend orientation?
no
Located near training center?
no
yes
Baseline surveys
yes
Highest value for training?
yes
no
Selected by owner?
yes
no
Potential Trainee
Non-Trainee
Treatment Firm?
Treatment Firm?
yes
no
yes
Trained Worker
Control: potential trained or spillover worker
Spillover Worker
Out of Sample

Table C5: Project Timeline

Oct 2022 - Jan 2023	**Listing and marketing**
	Introduced our program to the owners and all workers that present at the firm.
	Identified firms that would qualify to meet the number of employees collected, location, and contact information.
	Allowed firms an opportunity to apply to participate in our training program.
Feb 2023	**Orientaion**
	Our expert trainer introduced the curriculum to owners.
Jul - Nov 2023	**Baseline survey**
	Collected demographics from both owner and workers.
	Incentive-compatible exercises with the owner about perceived worker quality.
	Collected Owner selection for training.
	Collected worker demand for training, and thus, worker selection for training.
Aug - Nov 2023	**Baseline practical tests**
	Objective measure for worker quality at baseline for all workers in our evaluation sample.
Jan 2024	**Treatment assignment**
Feb-May 2024	**Training**
	Every cohort trains for two days a week and 6-8 hours per day.
May -June 2024	**Endline practical tests**
	Objective measure for worker quality at endline for all workers in our evaluation sample.

We invited workers from each location whenever we completed the baseline survey from that location as enumerators moved to another location during the baseline survey, which explains the overlap between baseline surveys and baseline practical tests.