

Online Appendix

The Uneven Impact of Generative AI on Entrepreneurial Performance: Evidence from a Field Experiment in Kenya

Nicholas Otis

Rowan Clarke

Berkeley Haas

Harvard Business School

Solène Delecourt

David Holtz

Rembrand Koning

Berkeley Haas

Columbia Business School

Harvard Business School

Table of Contents

A Appendix Figures	3
B Appendix Tables	19
C Details on the Experiment Design	30
C.1 Recruitment Strategy	30
C.2 Details on the Generative AI Assistant	31
C.3 Overview of the Experimental Interventions	32
D Detailed Case Examples	34
E Survey Questions and Measurement	48
E.1 Explanation of Preferred Functional Forms	49
E.2 Survey-based Mechanisms Questions and Analysis	49
F Ethical Considerations	54
G Experimental Spillovers	55
H Text Analysis	56
H.1 Overview and Summary Statistics	57
H.2 Random Forest Classification Details	58

H.3 Question and Answer Topic Analysis	59
H.4 Bag of Words Analysis	61
H.5 Treatment Effect Embedding Methodology	62
H.6 Implementation of Word Matching Analyses	63

A Appendix Figures

Figure A1: Distribution of study participants across Kenya

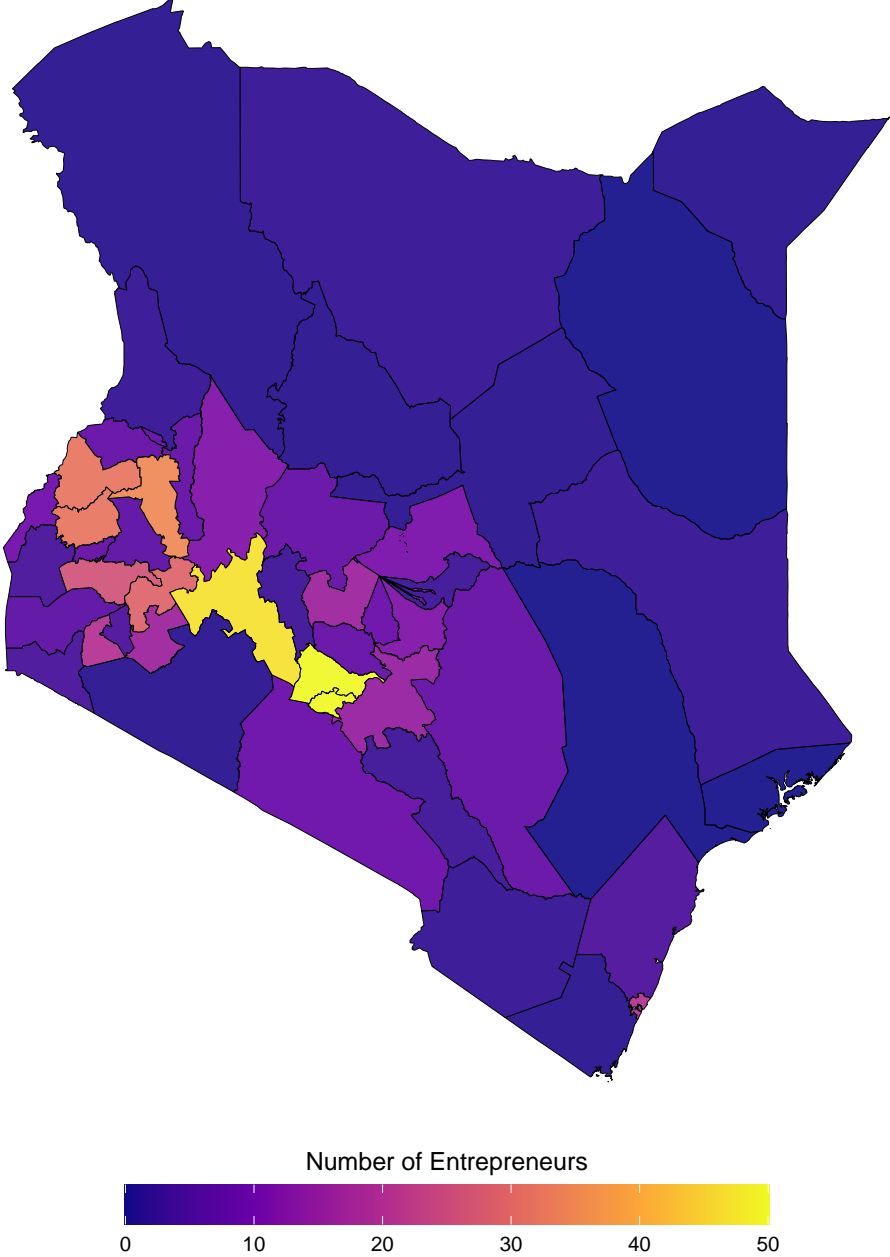
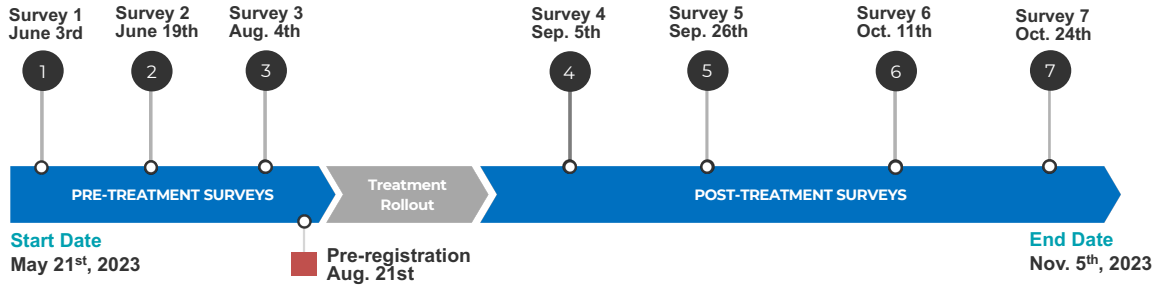
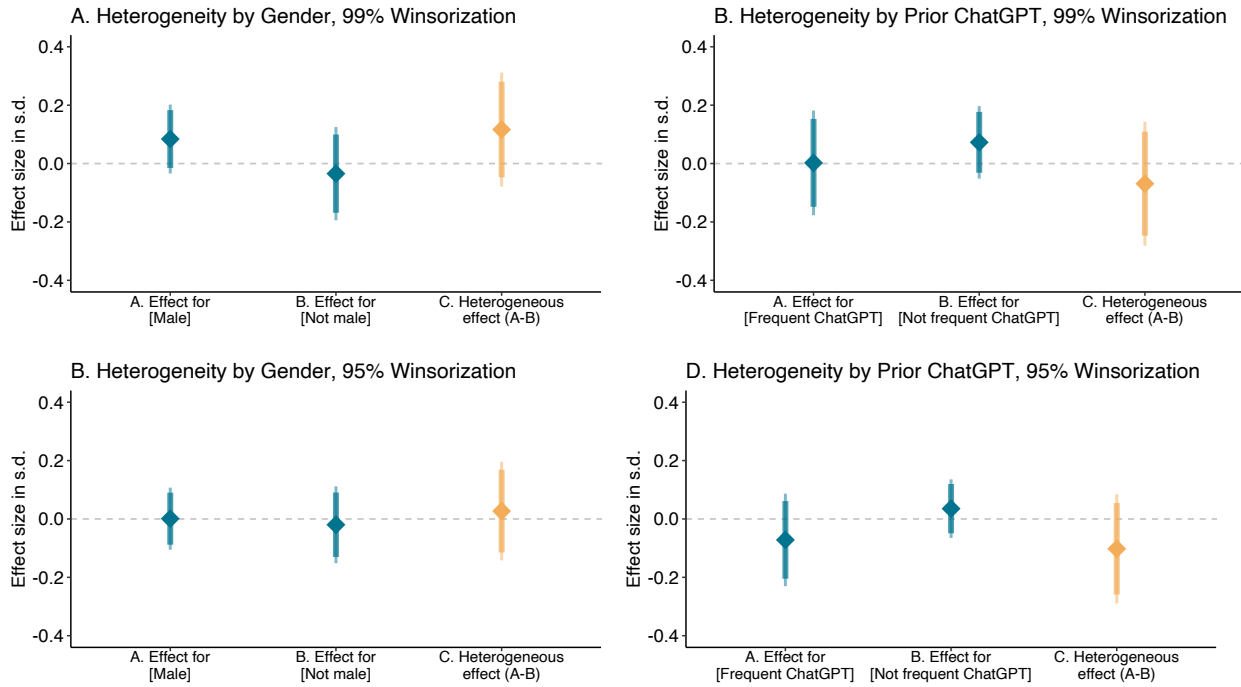


Figure A2: Overview of the experimental timeline



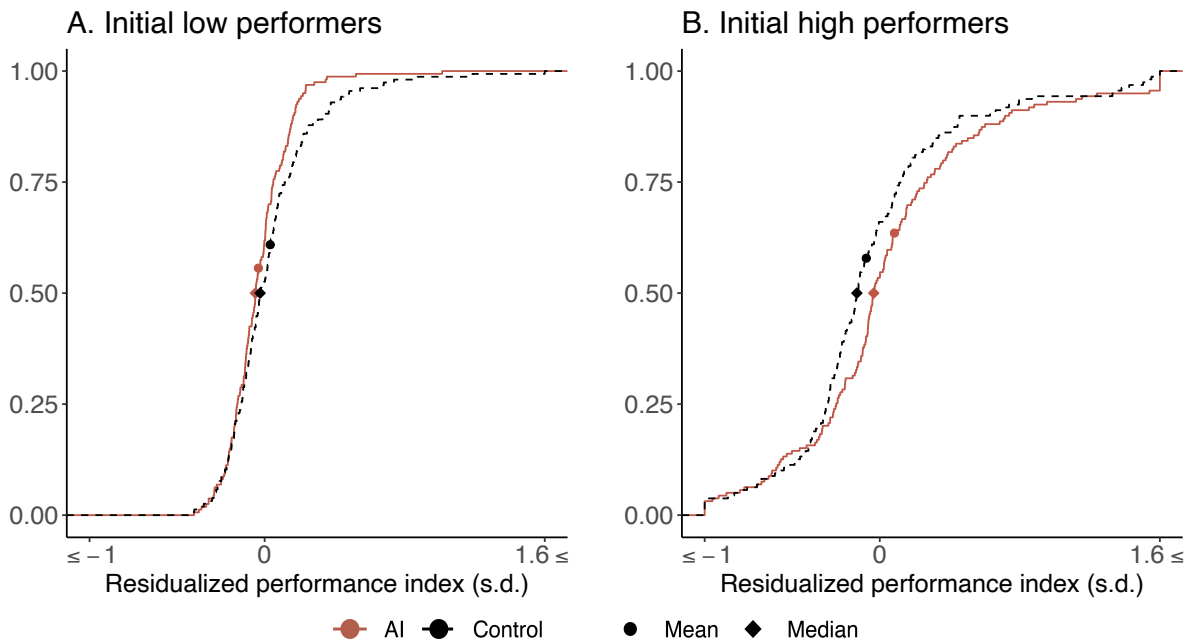
Overview of the experimental timeline. In total, there were three pre-treatment and four post-treatment performance surveys. The Start and End Date refer to the dates of the first and last surveys, while the seven dates presented above the timeline are the median survey completion dates for the seven performance surveys.

Figure A3: Performance effect heterogeneity by gender and prior ChatGPT use



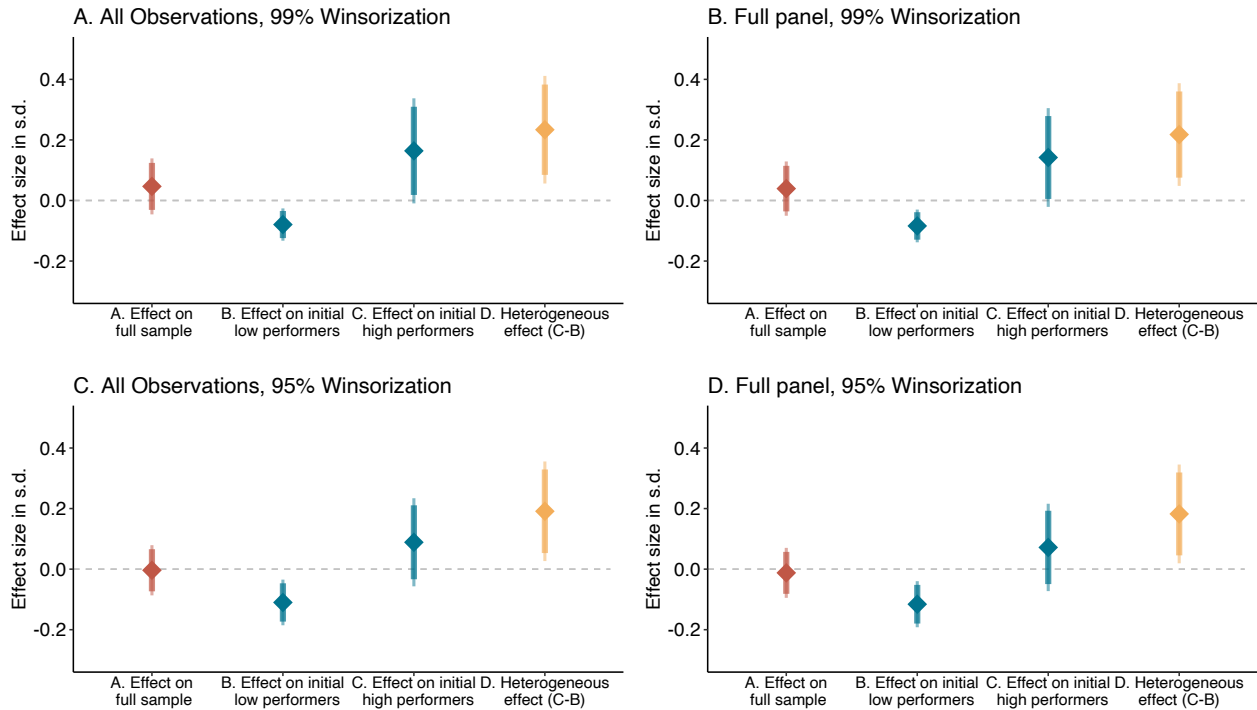
Heterogeneous effects of access to the AI assistant by gender and prior ChatGPT use. Panels A-B (C-D) display results winsorized at the 99% (95%) level. Panels A and C present heterogeneous effects based on participants' reported gender. Panels B and D report results based on prior ChatGPT use, where *frequent use* is defined as reporting using ChatGPT at least once per week. Effects are estimated controlling for covariates selected using double-LASSO, including pre-treatment business performance. Error bars represent 90% and 95% confidence intervals.

Figure A4: Empirical cumulative distribution functions of residualized business performance



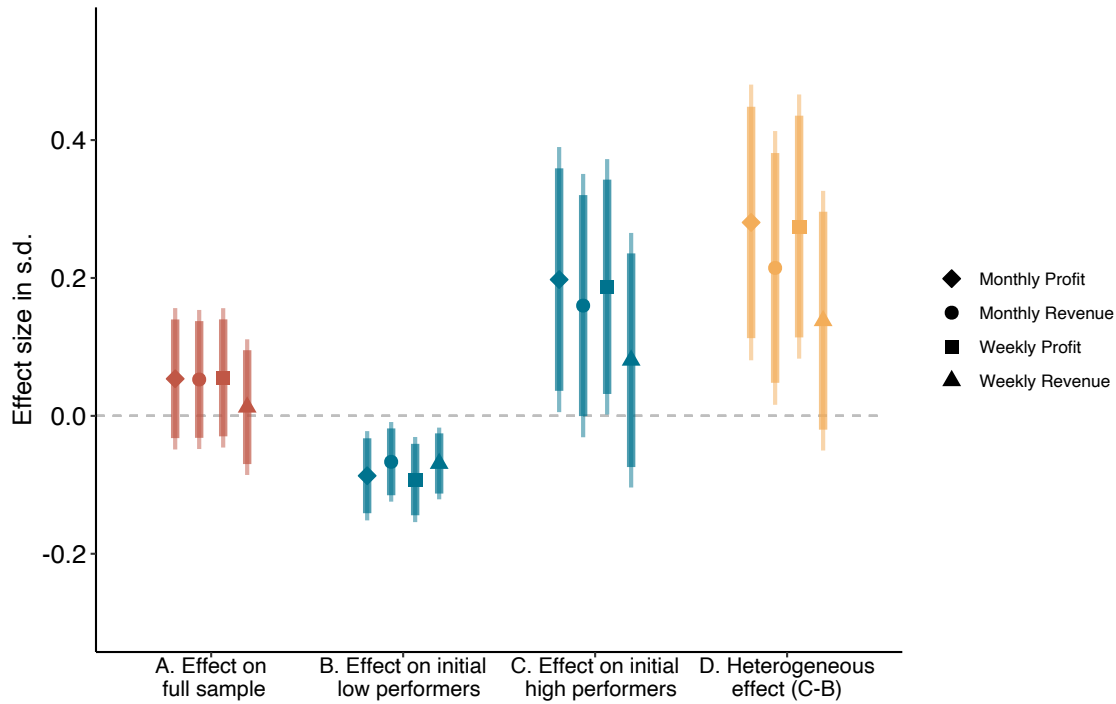
Panels A and B display empirical cumulative distribution functions (ECDFs) of the standardized business performance index. Red solid lines represent the AI assistant condition, and black dashed lines denote the control group. Panel A presents results for participants with below-median pre-treatment performance, while Panel B shows those with above-median pre-treatment performance. Higher values on the x -axis correspond to higher levels of business performance.

Figure A5: Tests of robustness to alternative sample inclusion criteria at both 95% and 99% winsorizations



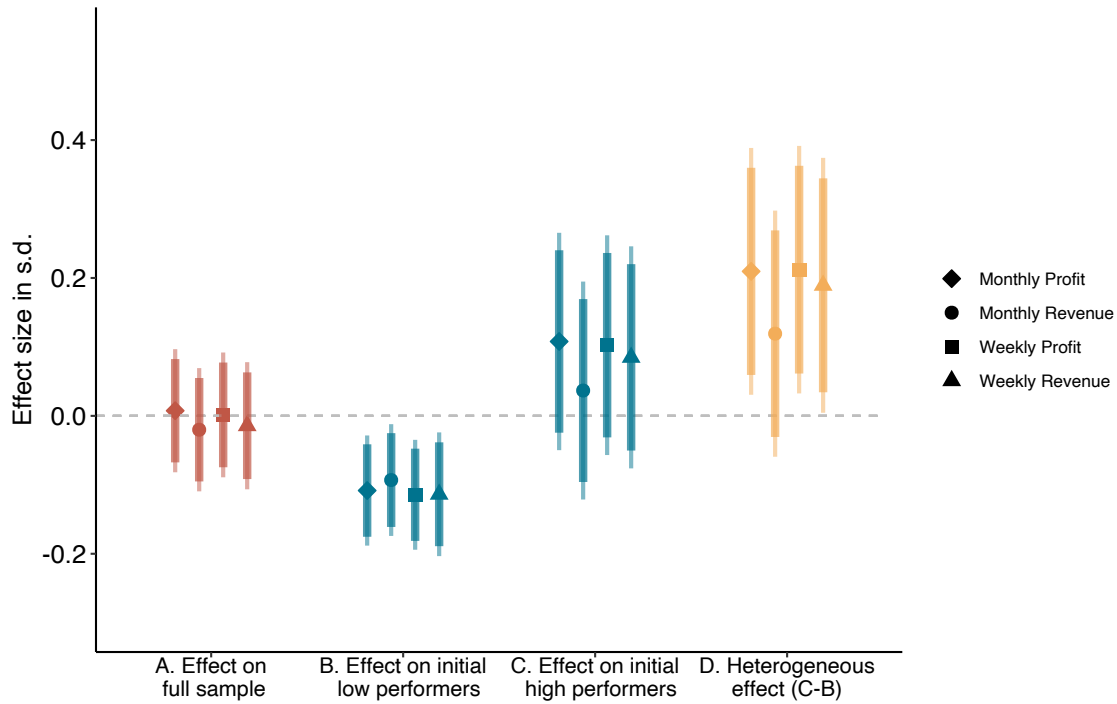
The effect of access to the AI assistant on the business performance index. Panels A and B (C and D) display results winsorized at the 99% (95%) level. Panels A and C include all observations, while Panels B and D restrict the analysis to participants completing all post-treatment surveys (i.e., “Full Panel”). In each panel, estimate A presents the average treatment effect of the AI assistant. Estimates B and C display results restricted to initially low- and high-performing firms, as defined by median value of the pre-treatment business performance index. Estimate D presents the heterogeneous treatment effect. Effects are estimated controlling for pre-treatment business performance, time fixed effects, strata fixed effects, and covariates selected using double-LASSO. Error bars present 90% and 95% confidence intervals.

Figure A6: Effects on individual performance index components winsorized at the 99% level



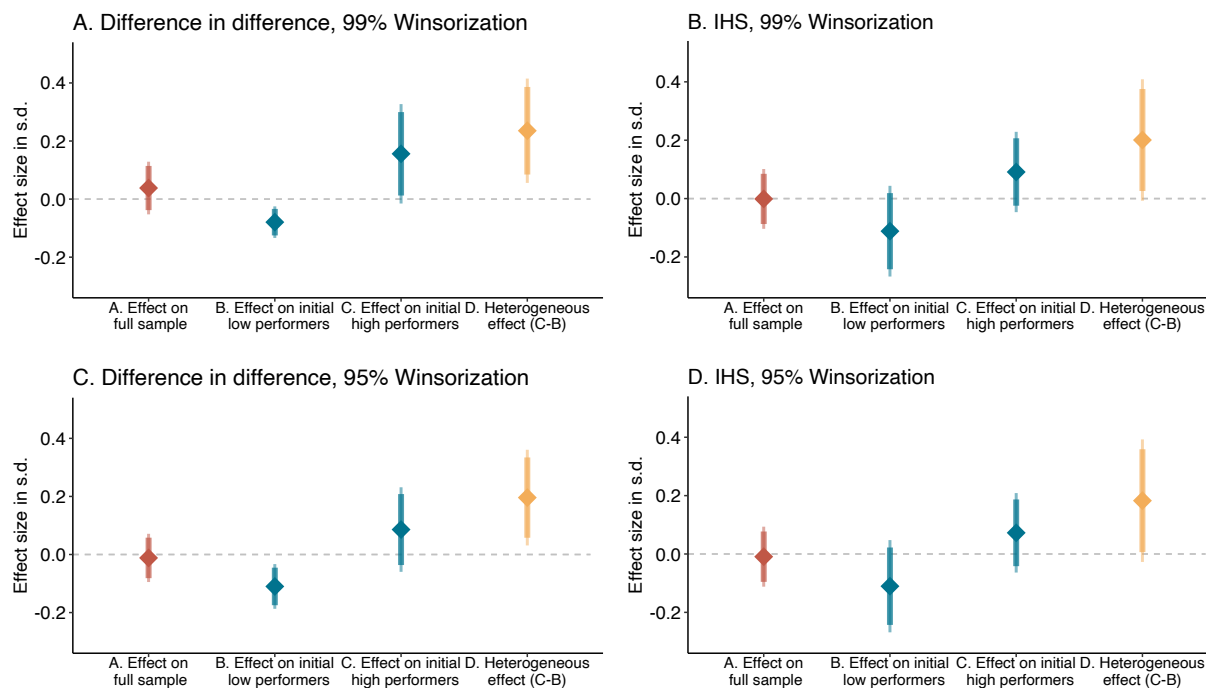
The effect of access to the AI assistant on the components of the business performance index (Weekly Revenue, Monthly Revenue, Weekly Profit, Monthly Profit). Results are winsorized at the 99% level. In each panel, estimate A presents the average treatment effect of the AI assistant. Estimates B and C display results restricted to initially low- and high-performing firms, as defined by median value of the pre-treatment business performance index. Estimate D presents the heterogeneous treatment effect. Effects are estimated controlling for pre-treatment business performance, time fixed effects, strata fixed effects, and covariates selected using double-LASSO. Error bars represent 90% and 95% confidence intervals.

Figure A7: Effects on individual performance index components winsorized at the 95% level



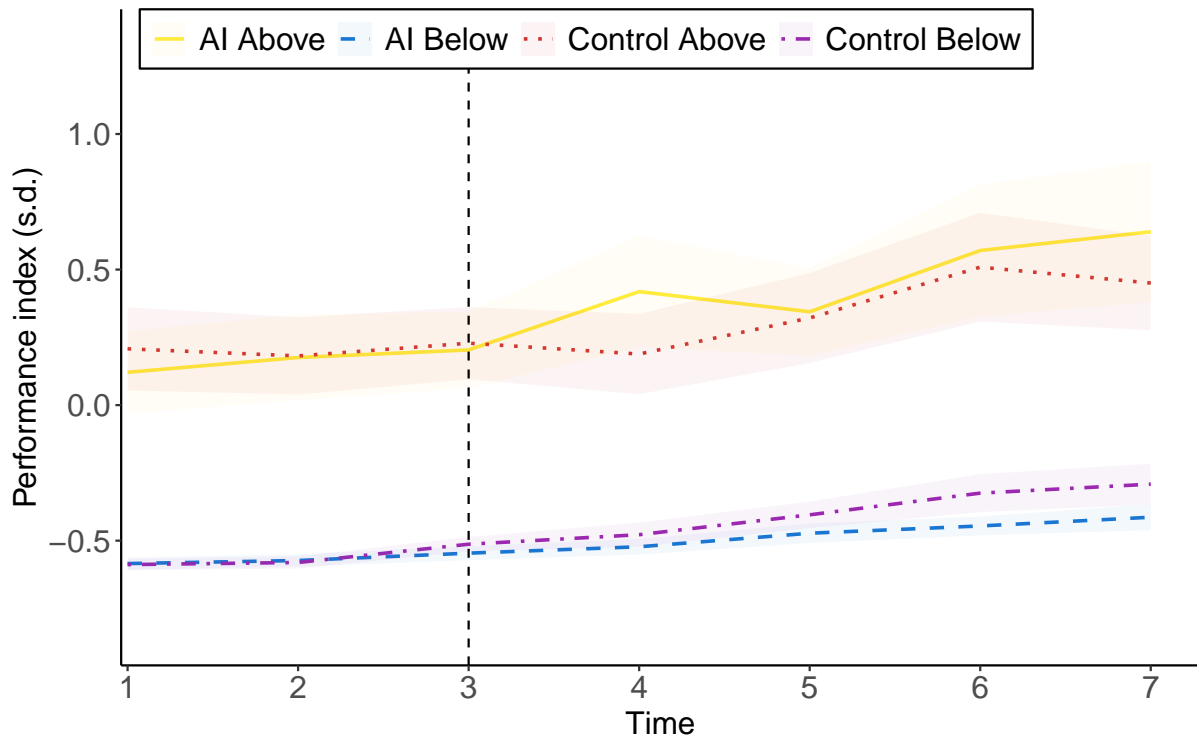
The effect of access to the AI assistant on the components of the business performance index (Weekly Revenue, Monthly Revenue, Weekly Profit, Monthly Profit). Results are winsorized at the 95% level. In each panel, estimate A presents the average treatment effect of the AI assistant. Estimates B and C display results restricted to initially low- and high-performing firms, as defined by median value of the pre-treatment business performance index. Estimate D presents the heterogeneous treatment effect. Effects are estimated controlling for pre-treatment business performance, time fixed effects, strata fixed effects, and covariates selected using double-LASSO. Error bars represent 90% and 95% confidence intervals.

Figure A8: Tests of robustness to alternative empirical specifications at both 95% and 99% winsorizations



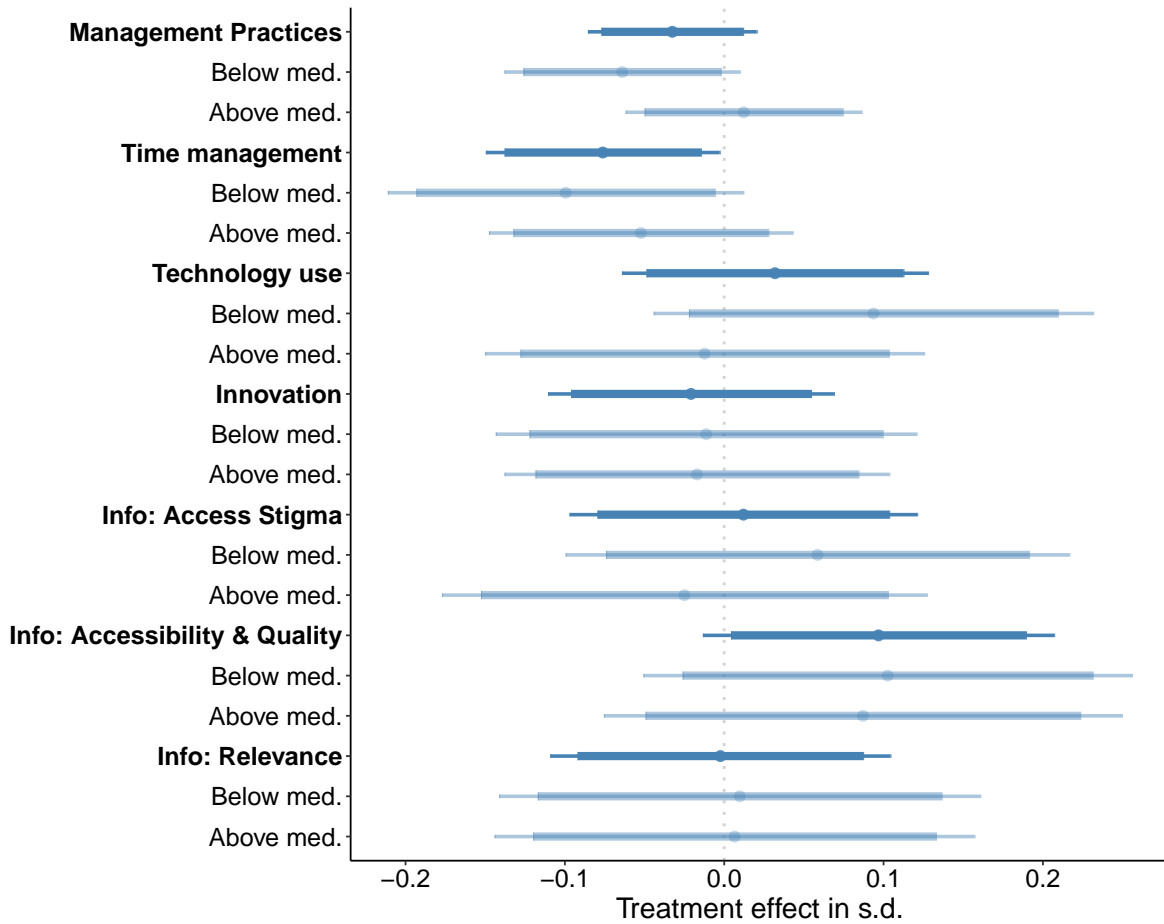
The effect of access to the AI assistant on the components of the business performance index (Weekly Revenue, Monthly Revenue, Weekly Profit, Monthly Profit). Panels A-B (C-D) display results winsorized at the 99% (95%) level. Panels A and C present results estimated using a difference-in-difference specification, and panels B and D present results for performance transformed using an inverse-hyperbolic sine (IHS) transformation. In each panel, estimate A presents the average treatment effect of the AI assistant. Estimates B and C display results restricted to initially low- and high-performing firms, as defined by median value of the pre-treatment business performance index. Estimate D presents the heterogeneous treatment effect. Effects are estimated controlling for pre-treatment business performance, time fixed effects, strata fixed effects, and covariates selected using double-LASSO. Error bars represent 90% and 95% confidence intervals.

Figure A9: Time trends in business performance index across subgroups



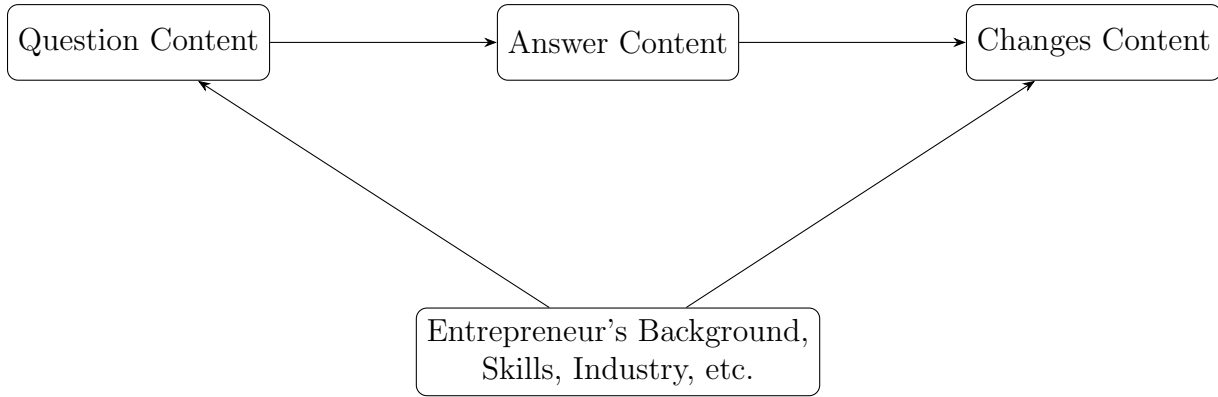
This figure displays time trends in the average performance index for entrepreneurs with above- and below-median pre-treatment performance who were randomly assigned access to the AI assistant. Access to the AI assistant was provided after the third time point. Colored bands present 95% confidence intervals.

Figure A10: Effects on survey-based mechanism questions



Treatment effect estimates for the seven preregistered mechanism indices. Results are presented for the average treatment effect (indicated by bold-titled index labels), as well as for participants with below- and above-median pre-treatment performance. Error bars represent 90% and 95% confidence intervals. Standard errors are clustered at the individual level.

Figure A11: Directed acyclic graph illustrating the backdoor criterion



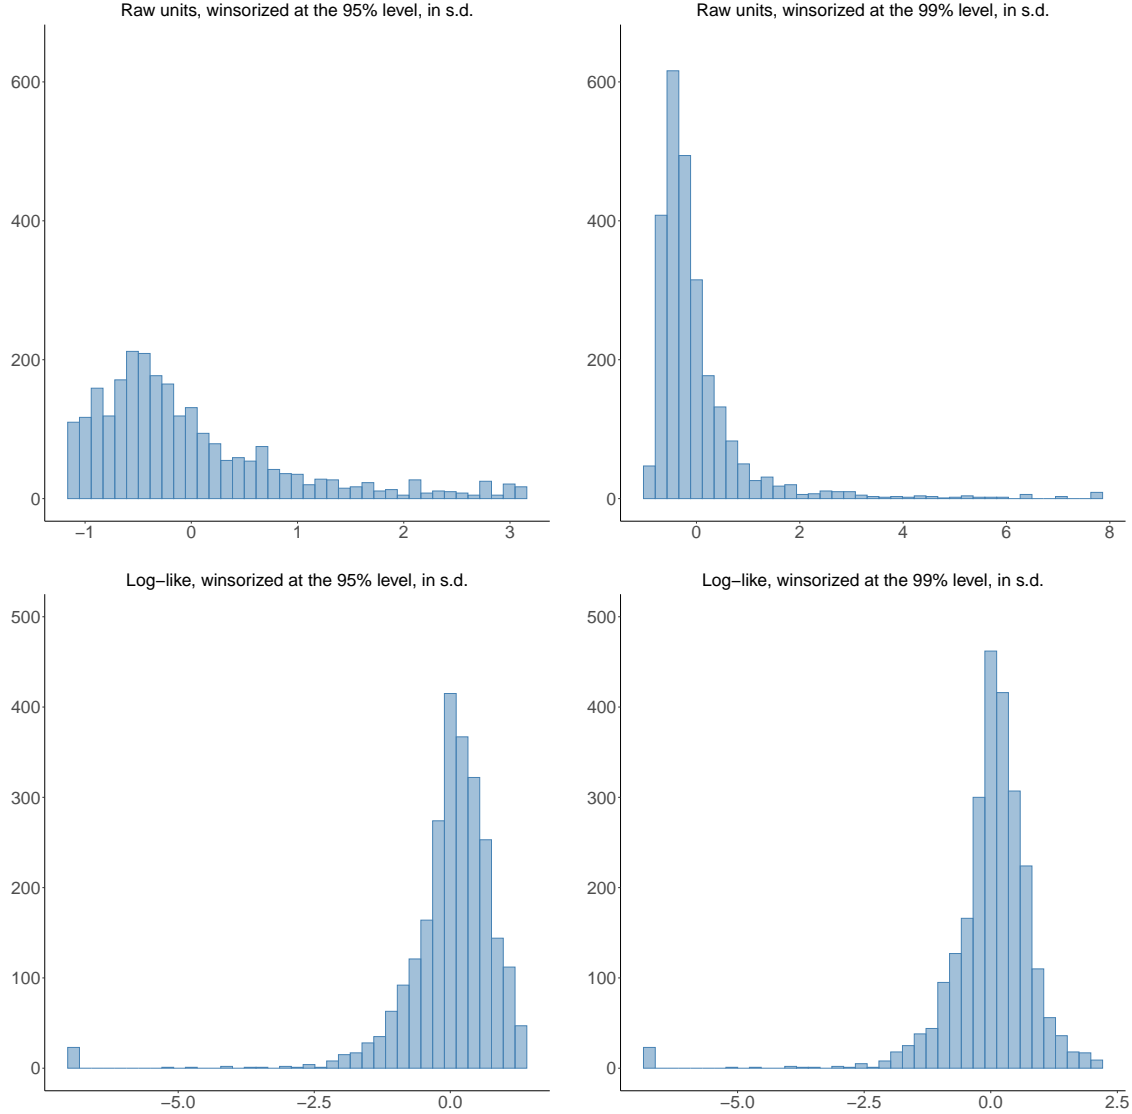
Entrepreneur characteristics influence both the questions asked and business changes made, creating potential confounding. Since the AI assistant only accesses information through WhatsApp messages, controlling for question content blocks all backdoor paths between AI advice and business changes, enabling causal identification of the AI's impact on entrepreneurial decisions.

Figure A12: Recruitment advertisement



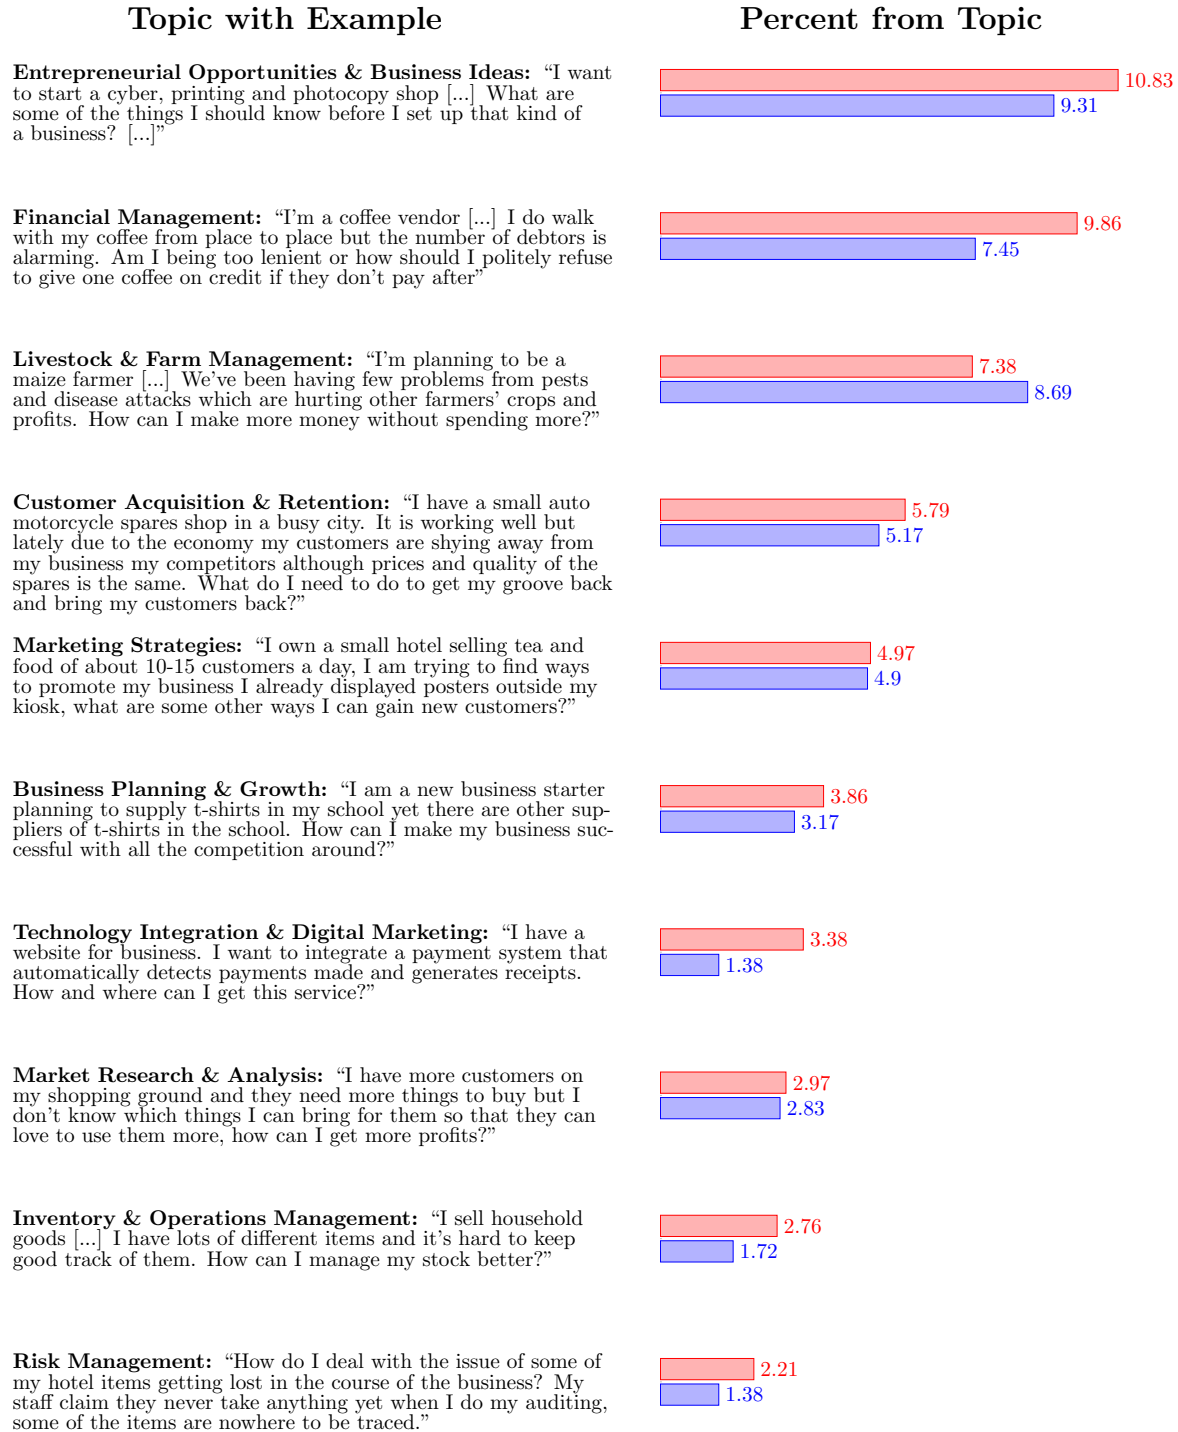
This figure presents the advertisement used to recruit participants to the study.

Figure A13: Distribution of performance index: Raw units vs. inverse hyperbolic sine transformation



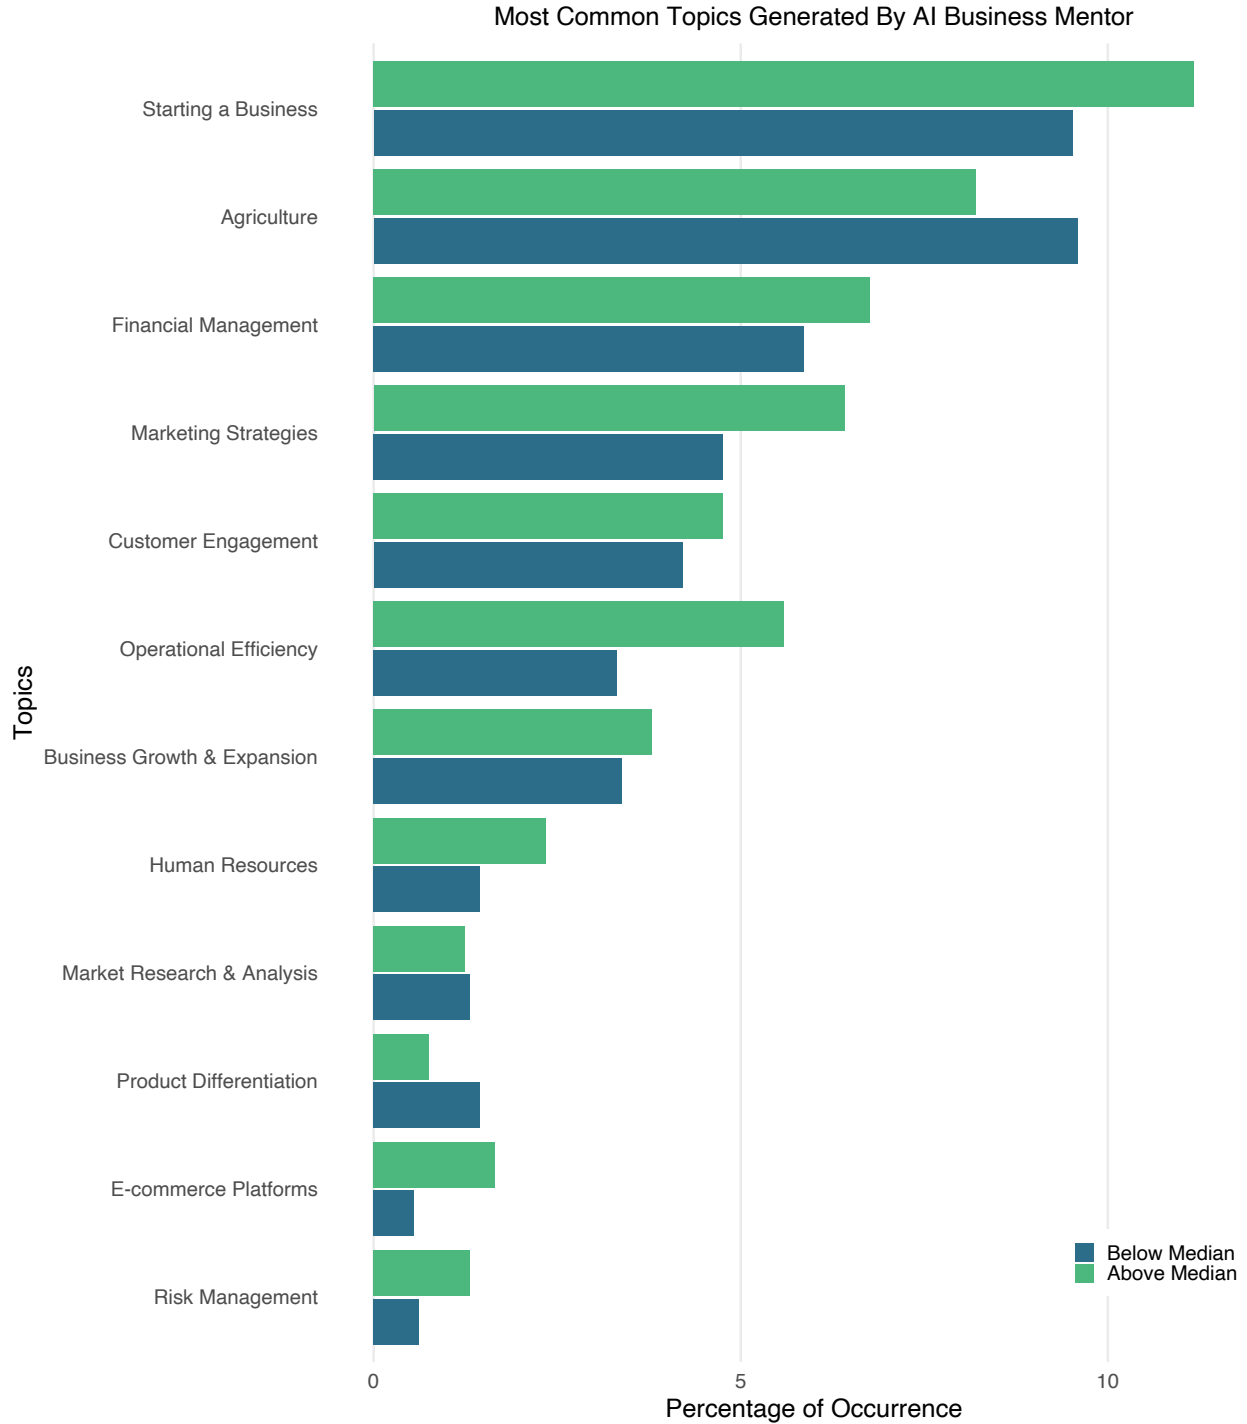
These figures present the distribution of four functional forms of our performance index. The top two figures depict results derived from raw units (Kenyan Shillings), and the bottom two depict results measured following an inverse hyperbolic sine transformation. In the first column results are winsorized at the 95% level and in the second column they are winsorized at the 99% level. In all cases, these results are presented in standard deviations for comparability. The bottom row illustrates why using raw units (via a “log-like” IHS transformation) is inappropriate, resulting in excess mass for revenue and profit values near zero. The top row shows the distribution of the performance index distribution constructed from “raw units” that we use throughout the paper.

Figure A14: Distribution of ten topics across questions sent to the AI assistant



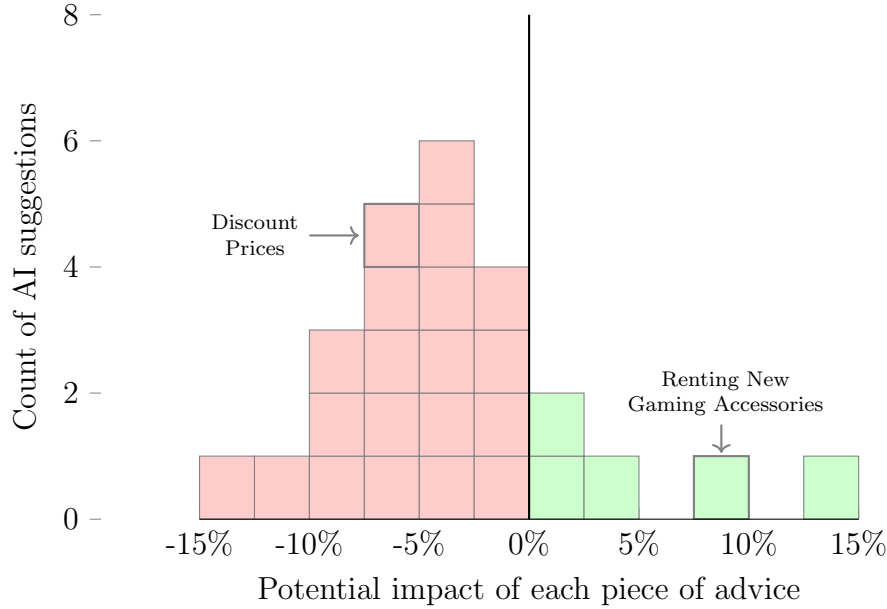
The bars show the percent of questions in each topic across initially low-performing entrepreneurs (red) and initially high-performing entrepreneurs (blue). See Appendix Subsection H.3 for additional details.

Figure A15: Distribution of Twelve Topics Across Responses Sent from the AI Assistant



The blue bars show the percentage of AI responses from each topic for below median low performers. The green bars show the percentage for above median high performers. See Appendix Subsection H.3 for additional details.

Figure A16: Illustration of Proposed Selection and Implementation Mechanism



This stylized figure illustrates how heterogeneous treatment effects can arise when entrepreneurs have different abilities to select and implement advice from a mixed-quality pool of AI suggestions. When some AI advice reduces performance (such as generic price-cutting strategies) while other suggestions meaningfully improve it (such as targeted product additions), variation in entrepreneurs' selection and implementation choices produces large outcome differences. Prior generative AI experiments have largely focused on well-defined tasks where suggestions are uniformly beneficial or harmful (e.g., Brynjolfsson, Li, and Raymond, 2025; ?), reducing the distribution of advice quality and minimizing the role of user judgment in determining outcomes.

B Appendix Tables

Table A1: Balance Table

	All	AI	Control	<i>p</i>
Business features				
Pre-treatment performance				
Monthly profit	19,550	19,010	20,097	0.52
Monthly revenue	40,959	41,134	40,781	0.92
Weekly profit	5,083	4,831	5,339	0.29
Weekly revenue	12,862	12,641	13,085	0.75
Other business features				
Management practices (count)	13.10	13.08	13.12	0.88
Oldest bus.: less than 1 year	0.20	0.20	0.21	0.86
Oldest bus.: 1 year	0.41	0.41	0.40	0.79
Oldest bus.: 2 years	0.23	0.22	0.24	0.71
Oldest bus.: 3 or more years	0.16	0.16	0.15	0.80
Sector				
Agriculture	0.19	0.18	0.19	0.78
Cosmetics	0.07	0.06	0.08	0.26
Carpentry	0.00	0.01	0.00	0.57
Information technology	0.11	0.14	0.09	0.05
Jewelry	0.02	0.02	0.01	0.21
Multimedia services	0.03	0.03	0.02	0.13
Clothing	0.21	0.22	0.20	0.49
Food	0.27	0.24	0.30	0.09
Tailor	0.03	0.03	0.03	0.84
Transportation	0.02	0.02	0.02	0.98
Construction	0.02	0.01	0.02	0.20
Other	0.08	0.09	0.08	0.60
Entrepreneur attributes				
Male	0.68	0.69	0.66	0.48
Children	0.73	0.71	0.76	0.48
Age	25.86	25.73	25.99	0.54
No college	0.10	0.09	0.11	0.56
Started college	0.37	0.37	0.37	0.90
Completed college	0.53	0.53	0.53	0.82
ChatGPT at least once a week	0.36	0.34	0.38	0.31
Agreeableness	9.57	9.46	9.68	0.09
Conscientiousness	10.21	10.15	10.27	0.28
Extraversion	8.45	8.31	8.60	0.07
Openness	10.20	10.13	10.27	0.23
Neuroticism	7.26	7.31	7.22	0.51
Location (county)				
Machakos	0.02	0.02	0.02	0.81
Bomet	0.03	0.03	0.02	0.63
Nyeri	0.03	0.02	0.03	0.98
Kisii	0.03	0.02	0.04	0.07
Mombasa	0.03	0.03	0.03	0.67
Kisumu	0.04	0.03	0.05	0.22
Kericho	0.05	0.04	0.05	0.82
Bungoma	0.05	0.06	0.04	0.49
Kakamega	0.05	0.06	0.04	0.49
Uasin Gishu	0.05	0.05	0.06	0.83
Nakuru	0.07	0.07	0.08	0.62
Kiambu	0.08	0.07	0.09	0.45
Nairobi	0.21	0.21	0.21	0.99
Other	0.26	0.28	0.24	0.24
Completed all surveys	0.97	0.97	0.98	0.35
Number of entrepreneurs	640	322	318	

Mean values and balance statistics for study participants. *p*-values capture a comparison across the different experimental conditions.

Table A2: Summary Statistics for Business Performance

	All	AI Assistant	Control
Panel A. Pre-treatment performance in Kenyan Shillings			
Mean: Monthly Profit	19,550	19,010	20,097
S.d.: Monthly Profit	23,420	23,807	23,021
Mean: Weekly Profit	5,083	4,831	5,339
S.d.: Weekly Profit	6,889	6,591	7,172
Mean: Monthly Revenue	40,959	41,134	40,781
S.d.: Monthly Revenue	47,038	48,598	45,428
Mean: Weekly Revenue	12,862	12,641	13,085
S.d.: Weekly Revenue	19,841	19,924	19,764
Panel B. Post-treatment performance in Kenyan Shillings			
Mean: Monthly Profit	25,844	25,987	25,699
S.d.: Monthly Profit	31,157	34,083	27,924
Mean: Weekly Profit	6,964	6,953	6,976
S.d.: Weekly Profit	9,365	9,958	8,733
Mean: Monthly Revenue	53,116	54,832	51,392
S.d.: Monthly Revenue	69,658	78,570	59,353
Mean: Weekly Revenue	15,833	15,702	15,965
S.d.: Weekly Revenue	22,567	23,459	21,642
Panel C. Main outcome variable: Standardized performance index			
Mean: Monthly Profit	0.01	0.01	0.00
S.d.: Monthly Profit	1.12	1.22	1.00
Mean: Weekly Profit	0.00	0.00	0.00
S.d.: Weekly Profit	1.07	1.14	1.00
Mean: Monthly Revenue	0.03	0.06	0.00
S.d.: Monthly Revenue	1.17	1.32	1.00
Mean: Weekly Revenue	-0.01	-0.01	0.00
S.d.: Weekly Revenue	1.04	1.08	1.00
Panel D. Performance surveys by group			
Total count of observations	4,434	2,226	2,208
Average observations per participant	6.93	6.91	6.94
Comparison of average counts by group (p-value)			0.41

Panel A reports the pre-treatment mean and standard deviation of the four business performance index components in Kenyan Shillings (monthly profits, weekly profits, monthly revenue, and weekly revenue), winsorized at the 99% level. Panel B reports the same statistics over the post-treatment period. Panel C presents standardized values for the main outcome, which by definition has a mean of 0 and a standard deviation of 1 for each index component in the post-treatment control group. Panel D reports on the total and average number of survey responses, which each response include values for all four of the business performance measures. The last row reports the p -value from a t -test comparing the average number of observations between the two experimental groups.

Table A3: Multiple Inference Correction by Families of Estimates

	p-values (1)	BH corrected p-values (2)
A) Heterogeneous treatment effects		
Pre-treatment performance	0.01	0.03
Gender	0.24	0.36
Frequent ChatGPT use	0.52	0.52
B) Survey-based mechanism effects		
Management practices	0.23	0.53
Time Management	0.04	0.30
Technology Use	0.52	0.90
Innovation	0.65	0.91
Info: Access stigma	0.83	0.96
Info: Accessibility and quality	0.09	0.30
Info: Relevance	0.96	0.96
C) Main performance effects		
Average treatment effect	0.32	0.32
Initial low performers	0.00	0.02
Initial high performers	0.07	0.09
Het. effect by pre-treatment performance	0.01	0.02

This table presents multiple inference-corrected p -values for three families of outcomes using the procedure outlined in Benjamini and Hochberg (1995) to control the false discovery rate. Column 1 shows the original p -values, and Column 2 shows p -values after applying the family-level correction. The three families are: heterogeneous treatment effects (Panel A), survey-based mechanism effects (Panel B), and main performance effects (Panel C).

Table A4: Top 100 Most Frequent Categories of AI Assistant Recommendations, with Number of Occurrences

Phrase	Count	Phrase	Count
Community Engagement	59	Vaccination	9
Market Research	51	Social Media Management	9
Poultry Farming	42	Soap Making Business	9
Loyalty Program	41	Second-hand Clothes Business (Mitumba)	9
Location	32	Risk Management	9
Inventory Management	31	Online Presence	9
Customer Feedback	29	Marketing & Sales Strategy	9
Value Addition	27	Market Demand	9
Pricing Strategy	26	Licenses and Permits	9
Leverage Social Media	26	Friends and Family	9
Partnerships	25	Food Business	9
Loyalty Programs	25	Fish Farming	9
Diversify Your Product Range	25	Feedback Mechanism	9
Crowdfunding	24	Executive Summary	9
Diversification	21	Emergency Fund	9
Customer Service	20	Affiliate Marketing	9
Offer Value-Added Services	19	Understand Your Customers	8
Marketing Strategy	19	Supplier Relationships	8
Local Partnerships	19	Rent	8
Improve Customer Service	19	Promotions & Discounts	8
Social Media Presence	18	Personal Savings	8
Social Media Marketing	18	Networking	8
Partnerships with Local Businesses	18	Location Selection	8
Budgeting	18	Laundry Services	8
Customer Loyalty Program	17	Healthcare	8
Referral Program	16	Freelance Services	8
Online Tutoring	16	Feeding	8
Business Plan	16	Education	8
Savings	15	Dropshipping	8
Improve Your Marketing	15	Customer Loyalty Programs	8
Agriculture	15	Tutoring Services	7
Optimize Operations	14	Training	7
Microfinance Institutions	14	Start Small	7
Crop Rotation	14	Soap Making	7
Create Customer Loyalty Programs	13	Seek Professional Help	7
Competition	13	Regular Communication	7
Promote Your Business Locally	12	Quality Control	7
Legal Requirements	12	Provide Excellent Customer Service	7
Negotiate with Suppliers	11	Promotion	7
Microfinance Institutions (MFIs)	11	Offer Promotions	7
Competitor Analysis	11	Offer Excellent Customer Service	7
Community Involvement	11	Marketing	7
Bulk Purchasing	11	Laundry Service	7
Record Keeping	10	Improve Product Display	7
Quality Assurance	10	Grants and Competitions	7
Promotions	10	Freelancing	7
Nutrition	10	Freelance Writing	7
Market Analysis	10	Educate Yourself	7
Email Marketing	10	Diversify Your Products	7
Word-of-Mouth Marketing	9	Content Creation	7

Table A5: Summary Statistics Describing the Messages Entrepreneurs Sent to the AI and the Messages the AI Sent Back

Panel A: Questions Sent							
(\geqMedian)	Entrepreneurs w/ Questions	Number of Messages Sent to the AI				Avg. Num.	Num. Words
		Average	S.D.	Median	Max	Followups	Per Question
FALSE	144	18.7	41	9	402	5.35	13.4
TRUE	137	15.5	19	8	105	4.56	14.5

Panel B: Answers Generated							
(\geqMedian)	Entrepreneurs w/ Answers	Pieces of Advice Generated				Num. Words	
		Average	Median	S.D.	Max	Per Answer	
FALSE	144	22.3	15	30.3	247	168	
TRUE	137	26.7	15	42.2	365	166	

Panel A reports the average number of messages sent by low- and high-performing entrepreneurs to the AI assistant, the average number of numeric follow-ups, and the average word count per question, excluding numeric follow-ups. Panel B presents information on the amount of AI-generated advice. Pieces of Advice Generated refers to the 3–5 suggestions outlined in our system prompt in response to non-follow-up questions classified as business-relevant (see Appendix H for details and Table A4 for the 100 most common advice categories).

Table A6: AI Advice and Entrepreneurs' Business Changes: Bag of Words

	In Changes Text					
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	0.18*** (0.01)	0.17*** (0.01)				
In AI Answer Text	4.15*** (0.39)	2.80*** (0.25)	2.99*** (0.25)	2.64*** (0.26)	2.64*** (0.26)	3.37*** (0.33)
In Question Text		14.50*** (1.86)	14.72*** (1.80)	13.61*** (2.31)	13.61*** (2.31)	15.78*** (1.99)
In AI Answer Text \times (\geq med.)				0.73* (0.33)		
In Question Text \times (\geq med.)				2.17 (2.43)		
Performers Subsample	All	All	All	All	Low	High
Entrepreneur			Yes	Yes	Yes	Yes
Observations	4,091,813	4,091,813	4,091,813	4,091,813	2,052,320	2,039,493

Words that appear in the AI's advice text are more likely to appear in the entrepreneur's description of their business changes. Figure A11 shows how controlling for the content of the question asked can allow us to recover the causal impact of the advice text on business changes. † $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Standard errors are clustered at the entrepreneur and word level.

Table A7: Regressions testing the heterogeneous impact of AI treatment on entrepreneurs' business changes

Panel A				
Dependent Variable:	Mentions AI Assistant?			
Model:	(1)	(2)	(3)	(4)
Reference group mean	0.63	0.64	0.64	0.63
AI	6.21***	1.97	1.84	10.33***
(\geq Median)	(1.47)	(1.42)	(1.40)	(2.54)
(\geq Median) \times AI		-0.39		
		(0.80)		
		8.51**		
		(2.92)		
Subsample	All	All	(<Median)	(\geq Median)
Observations	634	634	316	318

Panel B				
Dependent Variable:	Mentions Guides?			
Model:	(1)	(2)	(3)	(4)
Reference group mean	0.63	0.64	0.64	0.63
AI	-0.03	-0.59	-0.63	0.49
(\geq Median)	(0.62)	(0.64)	(0.63)	(1.03)
(\geq Median) \times AI		-0.18		
		(0.89)		
		1.13		
		(1.21)		
Subsample	All	All	(<Median)	(\geq Median)
Observations	634	634	316	318

Panel C				
Dependent Variable:	Discounts or Ads?			
Model:	(1)	(2)	(3)	(4)
Reference group mean	47.94	41.03	41.03	54.72
AI	3.71	11.93*	11.88*	-4.52
(\geq Median)	(3.59)	(5.20)	(5.19)	(4.93)
(\geq Median) \times AI		10.11 \dagger		
		(5.18)		
		-16.38*		
		(7.15)		
Subsample	All	All	(<Median)	(\geq Median)
Observations	634	634	316	318

Panels A and B test whether business change text mentions using the AI assistant or control guides via regular expressions. Panel C tests for mentions of discounting or advertising strategies. Reference group mean refers to the average value of the control group for the omitted category. Regressions control for the total number of words used. $\dagger p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Robust standard errors are presented in parentheses.

Table A8: Regressions testing the heterogeneous impact of using the AI assistant on the usage of uncommon words in the entrepreneur’s description of their business changes

Dependent Variable: Model:	Num. Uncommon Words			
	(1)	(2)	(3)	(4)
AI assistant	0.11 (0.07)	-0.03 (0.08)	0.00 (0.09)	0.19* (0.08)
(\geq med.)		-0.09 (0.10)		
(\geq med.) \times AI assistant		0.22† (0.13)		
Subsample	All	All	(< med.)	(\geq med.)
Observations	634	634	316	318

Poisson regressions showing high performers using AI employ more uncommon words when describing business changes. Examples of uncommon words include “detergent,” “headphones,” “peanut,” “rabbit,” and “uganda.” Common words include “advertising,” “customers,” “market,” “price,” and “vegetables.” Regressions control for the total number of words used. † $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Robust standard errors are presented in parentheses.

Table A9: Example Business Changes Illustrating Semantic Differences from the Treatment Effect Embedding for *High* Performers

Control Example	AI Treated Example
Similarity to high and low treatment effect embeddings: 3.7, 0.7	
Baking business. I've started baking honey filled muffins. Honey baked muffins. Home delivery of my products. Baking and supply of honey baked muffins.	Use of new technology of AI mentor. By doing more benchmarking on the competitors. Cooking oil, cereals. Use of AI mentor.
Similarity to high and low treatment effect embeddings: 3.5, 1.4	
Adding more products into the business. Delivering gas stove door to door. Oreo biscuits. Selling cakes in wholesale. Delivering. Refilling empty gas container. Delivering door to door. Advertising what am selling. The product is called bourbon. Delivering door to door services.	It has boost on sales and customer satisfaction. Improved on customer care skills after getting some tips and advice from AImentor and introduction of discounts to customers who purchased the products in bulk. These strategies have helped in customer loyalty. Introduction of groceries products into the business to reduce inconveniences to my clients. Yes, it has attracted new customers hence increase in sales and profit in the business. Introduced a new product line i.e selling groceries to reduce inconvenience to my clients.
Similarity to high and low treatment effect embeddings: 3.4, -0.4	
I have started going to the customer instead of waiting for them to come to me. I have started selling hair bonnets. I have started a delivery service. Delivery of products to customers has really made a difference in my business. Making myself available at any time for customers. I've started selling hair bonnets. Delivery.	AI teachings. Not Applicable. Saloon. Barbering. Use of AI to give me advice when needed. I have inserted a small salon beside my cyber to get more customers and profit. Use of solar when there's no electricity. Salon. A salon for women hair preparations. I got advice from AI mentor on other ways of getting power when there is electricity power black out. Having a salon in my barber shop and people being massaged after services. Selling magazines. Barbecue. Barbecue
Similarity to high and low treatment effect embeddings: 3.0, 0.6	
Adverting my products. Cooking quality food. Cooking chapati. Delivering my products to customers. Advertising my business through online. Forming Whatsapp group for my costumers. Selling Githeri. Giving my customers carrier bags. Selling ready cooked Githeri which has attracted many people. Delivering. Advertising my products online. Selling coffee. Delivering my products. Selling coffee	Improving the quality of chicks in my chicken roster has brought high quality chicken to my folder. I can now sell a chick that used to go for kshs 150 at 250 to 300 and this is a very great opportunity for me. I ma grateful for the Chat AI. I have used WhatsApp regularly to get information about market and business ideas. I have bought new chicks that were recommended by whe WhatsApp AI. I have regularly used bottles to drip feed my chicks. Selling of broilers and chick mash has been of great impact to my business. I don't rely on the chick and egg sales anymore, but get income from selling broilers and chick mash. Reducing slightly the prices of my commodities to help tilt the competition in my favour. I have started selling broilers and chick mash. Providing chick vaccination to farmers.
Similarity to high and low treatment effect embeddings: 2.7, -0.1	
Introduction of giids delivery. Initially, clients would pick their goods at the shop but now, they are delivered at their destinations t. Using Facebook ads to advertise the business. Initially i was selling cooking pots but i introduced other kitchen ware sets like cups and plates. Delivery of goods to customers. Engaging clients. Use of media in advertising. Male clothes. Delivery of products.	The incorporation of repairs in my shop has attracted more customers increasing both my income and sales. With the help of AI, I have been able to do more market research and discovered new ways of advertising my products. I have added cellphones, speakers and electric wires to my business. I have started electronics repairs e.g. Phones and Tvs alongside my business. The sale and refill of gas cylinders has really added a fortune to my operation. I have added the sale of Gas cylinders in my shop. I have started free delivery for Items worth 5000 and above.

Table A10: Example Business Changes Illustrating Semantic Differences from the Treatment Effect Embedding for *Low* Performers

Control Example	AI Treated Example
Similarity to high and low treatment effect embeddings: 0.9, 3.8	
Increase in profit in the business. Sales of eggs. Marketing my goods online. Selling of clothes. Transport service. Selling of clothes.	Giving so offers by reducing the selling price. To advertise in social media. Increase of sales. Advertising on social media platforms. Oraimo brand charger. Giving offers to some products.
Similarity to high and low treatment effect embeddings: 1.1, 3.3	
Online marketing. Getting to know what customers wants, precise record keeping and online marketing. Groceries. Cyber services. Reduced owing to reduce debts. Precise record keeping, balance sheet of the business. Cosmetics. Printing services. Marketing. Good marketing. Groceries. Providing loan at an interest.	Advertisement and discounts. Offering after sales service. Training of staff. Training of staff. Discounts. Offering discounts.
Similarity to high and low treatment effect embeddings: -0.5, 2.7	
Budgeting and data keeping. Budgeting and data keeping. Fruits. Mpesa services. Increased profits because I can now manage cashflow appropriately without possible losses. Planning and budgeting. Maize. Mpesa services. Fruit selling. I introduced selling of gapes ,oranges at my shop and it's adding up. Budgeting. Budgeting , planning , Recording. Mpesa services.	Posting online has brought more customers. A discount for items above ksh 1000. I started selling shoes in my clothes shop. Posting online my products. Afflicted customers. Afflicted customers. Mpesa alongside my shop. Discount for items above 1000. Offers for goods above 1500. Offers for goods above 1500. Selling shoes in my mtumba warehouse. Posting it on social media. Cereal outside my shop
Similarity to high and low treatment effect embeddings: 0.4, 2.5	
Consulting my customers. grocery. selling groceries. using of technology. selling of cosmetics products. Cosmetics. perfumes,. Grocery. Cosmetic shop. dealings with cosmetics products.	Isuing discounts on some products. Painting my shop as advertising strategy. Discount. Packaging.
Similarity to high and low treatment effect embeddings: -0.8, 2.4	
The mitumbas business. I have started selling materials for making vitenge. And including mitumbas. Inclusion of mitumba. To include selling shoes. Including mitumbas in my boutique. Bodaboda service. Including mitumbas in my business. Including mitumbas in shoes selling. Selling credit. Started layering poultry	The free delivery for my products is really working out for me. To work within the budget. I have tried the new enriched Amara oil. I have offered free delivery for my customers. Free delivery services to customers within range. I have offered free delivery for my clients. New Imperial leather for men. Hair dressing. It is a gaming house containing a play station 5 and Xbox. Free delivery services to my customers. Free delivery services. New saving gel. Free delivery services. Rearing and selling chicken

Table A11: Summary Statistics for the Aggregated Text Data

Panel A: Aggregated AI Questions

\geq Median	Entrepreneurs	Has Question(s)?	Num. Words Sent to the AI			
			Average	Std. Dev.	Median	Max
FALSE	160	0.900	109	170	56	958
TRUE	159	0.862	115	164	67	958

Panel B: Aggregated AI Generated Answers

\geq Median	Entrepreneurs	Has Answer(s)?	Num. AI Generated Words			
			Average	Std. Dev.	Median	Max
FALSE	160	0.900	2465	4069	1140	22854
TRUE	159	0.862	2197	3237	1018	19207

Panel C: Aggregated Business Changes

AI	\geq Median	Entrepreneurs	Has Text?	Num. Words to Describe Changes			
				Average	Std. Dev.	Median	Max
0	FALSE	156	0.821	46.1	44.5	35.5	204
0	TRUE	159	0.862	52.0	47.3	42.0	204
1	FALSE	160	0.850	44.9	43.1	33.0	182
1	TRUE	159	0.855	56.8	52.5	39.0	204

Panel A reports summary statistics for the questions that entrepreneurs ask the AI assistant. Panel B presents corresponding summary statistics for the responses from the AI assistant. Panel C reports statistics from the business changes reported by the entrepreneurs. In all panels, word counts are winsorized at the 99% level and are computed from a single entrepreneur-level text blob, constructed separately for each entrepreneur's set of questions, answers, and business changes.

C Details on the Experiment Design

This section provides additional details on the experiment design used in our study, including the recruitment strategy, the development of the AI assistant and its associated system prompt, and a short introduction to each experimental condition.

C.1 Recruitment Strategy

Our strategy for recruiting participants was designed to reach a diverse set of business owners throughout Kenya while minimizing the risk that expectations about the study would alter participants' behavior. Beginning in May 2023, participants were recruited through an advertisement run on the Meta ad platform in partnership with the Busara Center, a Kenyan research organization with more than a decade of experience.

The advertisement, presented in Figure A12, did not indicate that participants might later be offered access to either an AI tool or a business training guide. One potential advantage of this approach is that if participants knew in advance that some individuals would receive access to an AI assistant, they might be more inclined than non-participating individuals to independently access publicly available AI tools. Another potential benefit of this recruitment strategy is that it may have helped to avoid selection bias toward participants who might be, for example, especially interested in using AI. Note that it is still possible that participants who use platforms like Facebook were more likely than the average Kenyan to use AI. This could occur if, for example, Facebook users are more digitally literate than those who do not use Facebook or who did not click on the recruitment ad. However, it is worth noting that Kenya has high overall penetration (Kamau et al., 2024), meaning this recruitment strategy may be more representative in Kenya than if applied in other contexts.

The three pre-treatment surveys were used to identify and exclude respondents who stated that they did not own and operate a business (i.e., who were not SMB entrepreneurs). We also collected basic information on entrepreneurs' business performance, including their pre-treatment profits and revenue. Additionally, these three surveys helped to reduce attrition by excluding participants who failed to complete all three surveys or who failed basic attention checks. To encourage participation, participants received a small amount of compensation in

the form of airtime transfers for each subsequent survey that they completed. These transfers ranged in amount from 30 Kenyan Shillings for the first survey to 200 Kenyan Shillings for the four post-treatment surveys.

C.2 Details on the Generative AI Assistant

Our AI assistant was developed by integrating OpenAI’s GPT-4 API with the WhatsApp API. Participants were able to access the generative AI assistant using a public WhatsApp number. We configured the AI assistant so that it would only respond to users on a predefined access list, meaning that no participants in the control group or people outside of the study could interact with the system. To tailor the AI assistant to the Kenyan context, we developed a custom AI system prompt, without further model fine-tuning or adjustments to default tuning parameters.

We spent several months developing this AI system prompt. Initially, we worked with a team of research assistants to develop a prompt aimed at providing strategic advice to Kenyan entrepreneurs. Next, we solicited feedback from Kenyan entrepreneurs about the clarity, relevance, and practicality of the responses from the AI assistant, as well as on the format of the information it provided. Below is the final version of the system prompt used in the experiment:

As an expert business mentor, your specialization is guiding Kenyan entrepreneurs with limited resources. Your mission is to improve business performance and profitability by providing actionable, contextually relevant, specific, detailed advice that produces sustained improvements in competitive advantage. Interactions follow these steps:

- [1] When a question is asked, respond by learning about specific business needs relating to the business location, type, etc. Then, provide 3 to 5 actionable, easy-to-implement strategies tailored to the business in question. Offer a clear explanation for each strategy, including benefits & step-by-step guidance on implementation. Enumerate these options with [#].
- [2] After each message, prompt further discussion with the statement: “Enter a number to learn more about one of these points on *topic name* or ask a new question.” Offer a clear explanation for each tip, including the benefits & how to implement them.

- [3] If a number is given in the response, it indicates the desire to explore the corresponding strategy further. Offer 3 to 5 additional actionable tips related to the chosen strategy. Enumerate these options with [#].
- [4] Again, after each message, prompt further discussion with the statement: “Enter a number to learn more about one of these points on (bold topic name) or ask a new question.”

Communicate in the language of inquiry. Use simple non-technical English. Bold key advice and message topics. Avoid high-tech or costly solutions as your interlocutor is a Kenyan entrepreneur with limited resources. Avoid jargon. As a business mentor, do not discuss health, religion, politics, or current events. Only discuss business topics.

The most critical rule to remember: Whenever a number is entered, ALWAYS interpret it in reference to your last message. E.g., if the user responds “3”, it refers to point “3” in your most recent (last) message. If the user responds “2”, it refers to point “2” in your most recent message.

As described in the main text, the system prompt is designed to improve the quality of users’ interactions with the AI system by focusing the advice on the conditions of “Kenyan entrepreneurs” (as opposed to, for example, high-growth tech entrepreneurs), offering several enumerated suggestions in response to each query, and providing an easy interactive format that allows users to receive more details on any option by entering the number corresponding to that option.

C.3 Overview of the Experimental Interventions

Introduction to the AI Assistant

Because, at the time of our study, ChatGPT had only been introduced relatively recently, some Kenyan participants may not have had experience interacting with AI. As a result, in the condition that received access to the generative AI assistant, we started by providing a short overview of how the AI tool worked. We began by explaining that participants would be texting with an AI tool, and that the text responses would not be coming from a human.

We also explained that the AI tool might provide advice that is incorrect or unhelpful, and encouraged participants to think carefully about which advice they wanted to implement. Participants were also told that the AI assistant did not have up-to-date information—meaning, for example, participants could not ask it about yesterday’s weather.

As highlighted in Subsection C.2, the AI assistant was designed to respond to each query with enumerated responses. Furthermore, while the system prompt helped focus the AI’s advice on the Kenyan context, the assistant still often required substantial contextual information from the specific entrepreneur to provide meaningful advice. Participants were therefore advised to use two strategies to help them have productive interactions with the AI assistant: (1) providing detailed information about their situation, and (2) asking follow-up questions to the advice they received. For example, participants were told: *If you want to make more money as a farmer, don’t just ask “How can I make more money?” Instead, tell it about your farm, what crops you’re growing, and any problems you’re facing, like this: “I’m a sugarcane farmer in Nyando county with a 2-acre plot. We’ve been having a drought that’s hurting my crops and profits. How can I make more money without spending more?”*

Finally, we explained to participants how interactive conversations with the AI assistant were structured. This was necessary because the assistant was designed to provide a small number of enumerated suggestions in response to each query. In addition to being able to ask a new question at any time, participants could ask follow-up questions about any individual enumerated piece of advice.

Introduction to the ILO Guides

As noted in the manuscript body, the primary purpose of the PDF guides was to serve as a control condition that would have minimal impact on business performance. For this reason, the control group was simply provided with links to a series of business training guides from the ILO. At the time of the study, these guides were publicly available online, with past research suggesting that simply sending static training materials does not have any impact on business performance (Mehmood, 2024).

Participants were told that we were not affiliated with the International Labour Organization and were requested not to share the training materials they would be provided with. However, because these materials were at the time publicly available online, they could in theory be accessed by any participant.

Many participants may not have had experience navigating business guides before. As a result, we explained that each PDF started with a table of contents, which listed an

overview of what was covered in that specific guide. We also offered brief highlights from selected guides: participants were told that there were guides focused on coming up with new business ideas, starting a new business, or developing a business plan. Finally, participants were also told that the core ideas in the guides (e.g., conducting market research or developing sales projections) could help a business at any stage better understand the challenges it faces. As noted in Table A7, we find no evidence of user engagement with the guides, consistent with our interpretation of this group as a placebo control condition.

D Detailed Case Examples

Example 1: Dairy Products Business Owner (Below Median)

In this conversation, a milk seller sought advice on various aspects of their business. Initially, they faced a dilemma when a price hike from their supplier forced them to increase their retail prices, leading to the loss of some customers. They receive advice from the AI assistant on methods to maintain customer loyalty during this period, including transparent communication about the price increases, adding value to their products, ensuring quality, engaging with customers, and introducing flexible pricing options like smaller quantities and bulk discounts. The entrepreneur then expressed an interest in diversifying their product range to maximize profits. They were advised to consider selling items that complement milk, but faced challenges in finding reliable suppliers for these new products and received poor feedback from customers. In response, the AI assistant offered further suggestions, including online directories to find suppliers, attending trade shows, incentivizing customer feedback, and maintaining an open dialogue with customers to understand their preferences and improve product offerings.

Finally, the entrepreneur sought ideas to expand his business further after their attempt to sell yogurt was unsuccessful. They were advised to explore options like cheese making, starting a milk delivery service, selling milk-based snacks or desserts, expanding their wholesale business, and conducting training workshops on dairy products. Specific strategies for expanding into milk-based snacks and desserts were also discussed.

Q1. *Hello, I am [REDACTED]. I sell milk both wholesale and retail, in the past 1 week the price hiked from the supplier, so..I also decided to hike my price too in order to get some reasonable profit but now some customers have already left. What will I do to make them understand the prices are high from the production level?*

A1. Hello [REDACTED], it's a common challenge for businesses when costs increase. Here are some strategies to help your customers understand the situation:

1. **Transparent Communication**: Inform your customers about the increased prices from suppliers and explain that you had to adjust yours as well. This can be done through signage at your shop or via social media if you have an online presence. The key is to be open and honest about why prices have increased.
2. **Value Addition**: Enhance the perceived value of your milk products by offering something extra with each purchase such as loyalty points, discounts on bulk purchases, or small freebies like a packet of tea leaves or sugar.
3. **Quality Assurance**: Assure them that despite the price hike, the quality of your milk remains high. If possible, share information about how you ensure quality control in sourcing and storing your milk.
4. **Customer Engagement**: Engage more with your customers by asking for their understanding and support during these tough times while promising to review prices once things normalize at production level.
5. **Flexible Pricing Options**: Offer different pricing options where possible, such as smaller quantities at lower prices which may still be affordable for most of your customers.

Enter a number to learn more about one of these points on *customer retention strategies* or ask a new question.

Q2. 5

A2. Great choice, [REDACTED]. Implementing flexible pricing options can be a practical approach to retain customers during price hikes. Here are some additional tips on how you can do this:

1. **Small Quantity Pricing**: Offer smaller quantities of milk at lower prices. This could mean selling in half-liter or quarter-liter measurements for those who cannot afford a full liter due to the increased price.
2. **Bulk Discounts**: Provide discounts for customers buying in bulk. For example, if a customer buys 10 liters, they could get an extra liter free or at a reduced rate.
3. **Loyalty Program**: Establish a loyalty program where regular customers earn points with each purchase that can be redeemed for discounts or free products in the future.
4. **Payment Plans**: If feasible and applicable to your business model, consider offering payment plans for trusted regular customers who buy large quantities regularly but are struggling with the sudden increase in price.
5. **Promotional Offers**: Run occasional promotional offers such as 'buy one get one free' or 'price reductions on certain days' to keep customers interested and engaged despite the general price hike.

Enter a number to learn more about one of these points on *flexible pricing strategies* or ask a new question.

Q3. I'm [REDACTED], I sell milk both wholesale and retail but now, I want to add something else to the business in order to maximise my profit range. I don't know what will blend well with milk.

A3. Hello [REDACTED], it's great that you're thinking about diversifying your product range. Here are some suggestions of products that typically complement milk sales:

1. **Tea and Coffee**: These are commonly used with milk. You could consider selling various types of tea leaves or coffee beans.
2. **Cereal and Porridge**: Many people enjoy these with milk for breakfast. Offering a variety of cereals or porridge mixes could attract more customers.
3. **Bread and Pastries**: These items often go well with milk-based beverages for breakfast or snacks.
4. **Sweeteners**: Consider selling honey, sugar, or other sweeteners which are frequently added to milk.
5. **Dairy Products**: You can expand into other dairy products like cheese, yogurt, or butter which have a similar supply chain as milk.

Remember to research each potential addition thoroughly to understand the demand in your area, the cost implications, and how it might affect your operations.

Enter a number to learn more about one of these points on *product diversification* or ask a new question.

Q4. 5

A4. Excellent choice, [REDACTED]. Expanding into other dairy products can be a natural extension of your current business. Here are some actionable steps:

1. **Market Research**: Conduct surveys or informal chats with your customers to understand which dairy products they'd like to see in your store.
2. **Supplier Search**: Look for reliable suppliers who provide high-quality cheese, yogurt, and butter at competitive prices.
3. **Product Testing**: Initially introduce a small quantity of the new product to gauge customer response before fully stocking it.
4. **Promotion**: Inform your customers about the new products through signage, social media, or direct communication when they visit your shop.
5. **Feedback Loop**: Regularly ask for feedback from customers on these new products and adjust based on their responses - whether it's changing the brand, adjusting pricing or discontinuing if it doesn't meet expectations.

Remember that diversifying into dairy requires proper storage facilities due to perishability and quality control needs.

Enter a number to learn more about one of these points on *diversifying into other dairy products* or ask a new question.

Q5. 2

A5. Absolutely, finding the right suppliers is crucial when diversifying into other dairy products. Here are some steps to help you find reliable suppliers:

1. ****Identify Potential Suppliers****: Start by making a list of potential suppliers. You can do this by searching online, asking for recommendations from other business owners, or visiting local markets.
2. ****Evaluate Product Quality****: Request samples from these suppliers and assess their product quality. Remember, the quality of your products will reflect on your business reputation.
3. ****Compare Prices****: Ask for price lists from each supplier and compare them considering both the cost and quality of the products.
4. ****Check Reliability****: Find out about their delivery schedule, order minimums, return policy etc., to ensure they meet your specific needs.
5. ****Build Relationships****: Once you've chosen a supplier or two, build strong relationships with them as it can lead to better prices, more favorable terms and improved service over time.

Remember that a good supplier does more than just deliver products - they should be partners in helping your business succeed.

Enter a number to learn more about one of these points on **finding reliable dairy product suppliers** or ask a new question.

Q6. *Hello, this is [REDACTED] I sell milk both in wholesale and retail I contacted you asking what I can add in my shop that blends with milk and decided to explore other dairy products like cheese and yogurt now the problems that am facing are: (1). Supplier search i have not yet got a reliable supplier. (2). I have not been getting good feedback on why the customers are not loving the new products*

A6. Hello [REDACTED], I understand your concerns. Let's tackle them one at a time:

1. ****Reliable Supplier Search****: Here are some actionable steps you can take: [1] Use online directories or social media groups for Kenyan businesses to find potential suppliers. [2] Attend local trade shows or farmers markets where you can meet potential suppliers in person. [3] Ask other business owners in your network if they have recommendations for reliable dairy product suppliers.
2. ****Improving Customer Feedback****: [a] Conduct a brief survey with customers who've tried the new products. Ask specific questions about taste, packaging, price, and what they'd like to see improved. [b] Offer samples of the products to customers before they buy - this could help identify any issues early on and also encourage more sales. [c] Consider introducing smaller sizes or trial packs that cost less so that customers can try the products without feeling like it's a big investment.

Remember that changes in business often require patience and iterative improvements based on feedback. Don't be discouraged by initial setbacks.

Enter a number to learn more about one of these points on **supplier search** or **improving customer feedback** or ask a new question.

Q7. 2

A7. You've chosen to explore improving customer feedback further. Here are some additional steps:

1. **Open Dialogue**: Encourage customers to share their thoughts openly about the new products. Make them feel that their opinion is valued and important for your business improvement.
2. **Incentivize Feedback**: Offer small incentives like discounts or free samples in exchange for completing a feedback form.
3. **Direct Observation**: Watch how customers interact with the new products in your store. Do they seem confused? Are they comparing it with other items? This can give you clues about any potential issues.
4. **Trial Periods**: Run short trial periods for different products, then ask for feedback specifically on those items at the end of each period.
5. **Transparent Communication**: If you make changes based on customer feedback, let them know! This shows that you value their input and are committed to improving their experience.

Remember, *feedback is an opportunity* to learn more about your customers' needs and preferences, which can guide your business decisions moving forward.

Enter a number to learn more about one of these points on *improving customer feedback* or ask a new question.

Q8. *I sell milk both wholesale and retail in [REDACTED]. .now I wanted to know what else I can add to my business in order to maximize my profits. I tried yoghurts but didn't work.*

A8. Thank you for sharing about your business. It's great to see your interest in diversifying and increasing profits. Here are a few suggestions:

1. **Cheese Making**: With milk as your primary resource, cheese can be an excellent addition. Cheese has a longer shelf life than milk and can be sold at higher prices. Start with simple types like cottage cheese or mozzarella, which require minimal equipment and skills.
2. **Milk Delivery Service**: Consider offering home delivery services for regular customers within Nairobi. This service can increase customer loyalty, bring in consistent revenue, and allow you to reach more customers who value convenience.
3. **Selling Milk-Based Snacks or Desserts**: Depending on your location and customer base, adding snacks or desserts such as ice cream, milkshakes, or custards could attract more retail customers.
4. **Expand Your Wholesale Business**: Identify other businesses that might need dairy products (like restaurants or bakeries) and offer them competitive prices for bulk purchases of your milk.
5. **Training Workshops**: Conduct training workshops on the benefits of dairy products or how to make simple dairy-based foods at home (like homemade yogurt). This not only promotes your products but also creates an additional income stream from workshop fees.

Enter a number to learn more about one of these points on *business expansion* or ask a new question.

Q9. 3

A9. Great choice! Selling milk-based snacks or desserts can attract more customers and increase your profit margin. Here are some steps to implement this strategy:

1. **Market Research**: Start by understanding what kind of milk-based snacks or desserts are popular in your area. You could do a quick survey with your current customers or observe what's selling well in nearby food businesses.
2. **Simple Recipes**: Choose simple recipes that require minimal additional ingredients and equipment. Puddings, custards, or simple ice creams could be good starting points.
3. **Pricing**: Price these items competitively while ensuring you still make a profit after considering the cost of ingredients, preparation time, and utilities.
4. **Promotion**: Use signage at your shop to promote these new products and consider offering samples to regular customers to get their feedback and generate interest.
5. **Quality Control**: Ensure the quality of your desserts is consistently high so that customers keep coming back for more.

Enter a number to learn more about one of these points on *milk-based snacks/desserts* or ask a new question.

Q10. 1

A10. Absolutely, conducting market research is a crucial first step. Here are some actionable steps:

1. **Surveys**: Create simple surveys asking your customers about their favorite milk-based snacks or desserts and if they'd be interested in buying them from your shop. This could be done informally when they visit your shop or through a more structured approach like questionnaires.
2. **Observation**: Visit other food businesses in your area that sell milk-based snacks or desserts. Note which items seem popular, how they're priced, and any unique selling points.
3. **Community Engagement**: Engage with local community groups or online forums to understand what types of desserts people in Nairobi enjoy and are currently trending.
4. **Trial Products**: Once you've gathered this information, select one or two products that seem promising and start offering them on a trial basis to gauge customer response.
5. **Feedback Collection**: Collect feedback from the customers who try these new products - did they enjoy it? Would they buy it again? How much would they be willing to pay for it?

Remember: The goal of market research is not just to identify what products you should offer but also to understand how you can differentiate yourself from competitors.

Enter a number to learn more about one of these points on *market research* or ask a new question.

Example 2: Fast Food Joint Owner (Above Median)

In this conversation, the owner of a fast food joint located near a matatu terminus sought advice on differentiating their business in a competitive environment. The key strategies discussed included differentiating the menu with unique or special dishes, improving customer service, implementing a loyalty program, engaging in community activities, and maintaining high standards of cleanliness and hygiene.

The discussion delved into the specifics of establishing a loyalty program, with the AI assistant offering suggestions like a point-based system, special offers for regular customers, and birthday or anniversary rewards. The entrepreneur raised concerns about the cost implications of differentiating the menu and the complexities of community involvement. The AI assistant then suggested using affordable local ingredients, offering smaller portions, running specials during off-peak hours, using simple loyalty systems like punch cards, and engaging in low-cost community activities.

The entrepreneur was also concerned about the risks and customer reception related to changing the menu and implementing new strategies. The AI assistant recommended testing new dishes gradually and seeking customer feedback, especially from regulars. The importance of using both in-person and digital channels to collect reliable feedback was emphasized, considering that some customers are transient. The entrepreneur and the AI assistant also discussed the risks associated with implementing new business strategies, with the AI assistant suggesting thorough market research, pilot testing, financial analysis, risk assessment, and regular feedback review.

Q1. *I have fast food joint that I run just adjacent to a matatu terminus. There are several such food joints at that particular place and it's fair to say competition is stiff. How can I overcome the competitive business environment and be ahead of the rest of the pack?*

A1. Thank you for sharing about your business. Operating a fast food joint in a highly competitive environment like a matatu terminus can indeed be challenging, but there are several strategies you could consider to stand out:

1. **Differentiate Your Menu**: Offer unique or special dishes that aren't available at the other joints. This could be traditional Kenyan foods with a twist, or it could involve using locally sourced ingredients that others don't use. The benefit is that customers will associate your place with these special offerings and choose you over competitors.
2. **Improve Customer Service**: Prioritize excellent customer service to make people feel welcome and appreciated when they visit your joint. This might involve training staff on how to interact with customers positively, quickly serving orders, and handling any complaints promptly and professionally.
3. **Loyalty Program**: Implement a loyalty program where customers earn points for every purchase which they can later redeem for free meals or discounts. It encourages repeat business as customers will want to accumulate more points.
4. **Community Involvement**: Engage in local community events or sponsor local initiatives such as sports teams or school activities. This raises your profile within the community and makes people more likely to support your business.
5. **Cleanliness & Hygiene**: Ensure high standards of cleanliness both inside the kitchen and in the dining area because hygiene matters greatly in food businesses.

Enter a number to learn more about one of these points on *improving competitiveness* or ask a new question.

Q2. 3

A2. Great, let's delve into how you can implement a **Loyalty Program** in your fast food joint:

1. **Point-Based System**: Implement a simple point-based system where customers earn points for every shilling spent. Once they accumulate a certain number of points, they can redeem them for free meals or discounts. You could use simple tools like punch cards or digital apps that are inexpensive to manage.
2. **Special Offers for Regular Customers**: Identify your regular customers and offer them exclusive deals or discounts. This not only rewards their loyalty but also encourages them to keep coming back.
3. **Birthday/Anniversary Rewards**: Offer special treats or discounts to customers on their birthdays or anniversaries if they share this information with you. This personalized touch can make customers feel valued and appreciated.
4. **Referral Program**: Encourage your loyal customers to bring in new clients by giving them incentives such as a discount or free meal when someone they refer makes their first purchase.
5. **Feedback Reward**: Encourage customer feedback by offering small rewards for completing surveys about the quality of food and service at your joint.

Remember, the goal of any loyalty program is to make your customers feel valued so that they continue patronizing your business over others.

Enter a number to learn more about one of these points on **loyalty programs** or ask a new question.

- Q3. 1. **Differentiating my menu** means that I'm incurring a higher cost of preparing meals. This in turn pushes the menu prices upwards which may scare away the potential customers considering the fact that the potential customers are low or middle income earners. 2. About **loyalty programs**, I'm not well conversant with that field. 3. **Community involvement** may require adequate resources and may be capital intensive.

A3. Thank you for your detailed response. Let's delve into these topics a bit more:

1. ****Differentiating Your Menu****: You're correct that introducing unique dishes may increase costs. However, this doesn't always have to result in higher menu prices. Here are some tips: [1] Focus on dishes that use affordable local ingredients but are prepared in a distinctive way. [2] Consider offering smaller portions at lower prices for those who want to try something new without paying a lot. [3] Run specials during off-peak hours to attract customers and offset the cost of unique dishes.

2. ****Loyalty Program****: A loyalty program is essentially a reward system for frequent customers. Here's how you could implement one with minimal costs: [1] Use simple paper punch cards where each purchase gets a punch, and after a certain number of punches, the customer gets a free meal or discount. [2] Keep track of repeat customers manually and offer them discounts or free items after they've made several purchases.

3. ****Community Involvement****: While it can be resource-intensive, there are ways to engage with the community without spending much money: [1] Offer your venue as a meeting place for local groups during slow business hours. [2] Sponsor small-scale local events like neighborhood cleanups by providing refreshments - this helps build goodwill within the community.

Enter a number to learn more about one of these points on **cost-effective strategies** or ask a new question.

- Q4. *If we focus on having a menu that stands out from what others are offering, by what margin may I lose, retain or increase my customer base. And what is the likelihood of having rewards for customer's referrals?*

A4. The impact of differentiating your menu on your customer base can be significant, but it's hard to quantify without specific data. However, here are some general insights:

1. ****Potential Increase in Customers****: Offering unique dishes can attract food lovers who are always looking for something new and exciting to try. If your new offerings resonate with the taste preferences of the local population, you could see a substantial increase in customers.

2. ****Retention of Existing Customers****: If you continue offering popular items while introducing new ones, you're likely to retain most of your existing customers. They might even appreciate the added variety.

3. ****Possible Losses****: There is always a risk that some customers may not like changes to the menu and choose other joints over yours. It's essential to introduce changes gradually and get feedback from customers before making any drastic alterations.

As for referral rewards:

4. Referral programs can be highly effective because people trust recommendations from friends or family more than advertising messages from businesses. If executed well, this strategy could lead to a significant increase in new customers without requiring large investments in marketing.

To implement these strategies:

A) For **Menu Differentiation**, start by researching what foods or flavors are trending locally or globally that align with your capabilities and resources. Test these out as specials before adding them permanently to your menu based on customer response.

B) For **Referral Rewards**, consider giving both the referring customer and the referred friend a discount or free item on their next purchase. Make sure it's easy for customers to share referrals by providing them with simple instructions when they make purchases.

Enter a number to learn more about one of these points on **menu differentiation** or **referral rewards** or ask a new question.

Q5. *I'll need to contemplate on this two strategies and evaluate the outcomes both on a short and long term basis.*

A5. Absolutely, taking the time to contemplate and evaluate potential outcomes is a crucial step in making strategic decisions for your business. Remember that both **menu differentiation** and **referral rewards** are strategies aimed at improving customer retention and attracting new customers, which are key aspects of business growth.

When evaluating these strategies:

1. Consider your resources: Do you have the capacity to introduce new menu items or manage a referral program effectively?

2. Listen to your customers: What do they want or need? Would they appreciate a more diverse menu or would incentives for referrals motivate them?

3. Monitor results: Once implemented, track the success of these strategies through increased sales, more repeat customers, or an increase in new customers.

4. Be ready to adjust: If something isn't working as well as expected, don't be afraid to make changes.

Remember that successful businesses are those that adapt and evolve based on their environment and customer needs.

Feel free to return if you have any other questions or need further clarification on any point related to improving your fast food joint's competitive advantage

Q6. *Customer feedback when venturing into new ideas is the key to monitor the reception of a new menu and also the trust built through referral programs may work well with the business and the projected profit margins.*

A6. Absolutely, you're on the right track. Customer feedback is invaluable when introducing new ideas like a different menu or a referral program. Here's how you can effectively use customer feedback and trust building:

1. ****Feedback Channels****: Establish easy ways for customers to give feedback such as suggestion boxes at your joint, conducting informal surveys, or simply talking to them about their experience.

2. **Respond to Feedback**: Show customers that you value their opinions by acting on their suggestions where feasible and communicating changes back to them.
3. **Trust Building in Referral Programs**: Make sure the process is transparent and rewards are delivered as promised. This will build trust not only with the referring customer but also with the new ones coming in through referrals.
4. **Evaluate Profit Margins**: Regularly review your costs versus revenues after implementing these strategies to ensure they're contributing positively towards your profit margins.
5. **Pilot Test**: Before rolling out any major changes, consider doing a small-scale pilot test to gauge reception and adjust based on feedback received.

Remember, open communication with your customers helps foster loyalty which can lead to increased profitability in the long run.

Enter a number to learn more about one of these points on *customer feedback* and *trust building* or ask a new question.

Q7. *Together with my team, we shall deliberate on few important tips I've come across here and decide the way forward in best way to implement them.*

A7. That sounds like a great plan. Collaborative decision-making often leads to more effective strategies and smoother implementation. Here are some tips for your deliberation process:

1. **Set Clear Goals**: Clearly define what you want to achieve with each strategy, such as increasing customer retention, attracting new customers, or boosting profit margins.
2. **Assign Roles**: Assign team members specific roles in the implementation of each strategy based on their skills and responsibilities.
3. **Plan for Feedback**: Plan how you will collect and analyze customer feedback once the changes are implemented.
4. **Prepare for Adjustments**: Be prepared to make adjustments based on the outcomes of your initial efforts.
5. **Schedule Follow-ups**: Schedule regular follow-up meetings to assess progress and make necessary adjustments.

Remember that it's important to have everyone on board with these changes so they can be effectively implemented and managed over time.

Enter a number to learn more about one of these points on *team deliberation* or ask a new question.

Q8. *Good insight. I'll come back here sooner with some more insights on the progress made if any.*

A8. That's great to hear! Remember, the process of implementing new strategies often involves learning and adjusting along the way. Don't hesitate to return if you have any questions or need further advice as you progress. Wishing you and your team all the best in enhancing your business performance and profitability.

Q9. *Hi... If I'm running a food joint near a matatu terminus, selling fast food, snacks and beverages, how can I diversify the services I offer*

A9. Hello! Diversifying your services can be a great way to attract more customers and increase revenue. Here are some ideas:

1. ****Delivery Service****: You could start offering delivery service for customers who may not have time to visit your food joint in person but want a quick meal.
2. ****Catering Services****: Consider providing catering services for local events or businesses, which can significantly boost sales.
3. ****Meal Deals****: Introduce special meal deals during specific hours of the day, such as breakfast combos or late-night snacks.
4. ****Seasonal Menu Items****: Offering seasonal menu items that align with holidays or local events can attract new customers and keep regulars interested.
5. ****Loyalty Program****: Implement a loyalty program where regular customers earn points they can redeem for meals or discounts.

Each strategy has its own benefits and methods of implementation, so consider what would work best for your specific business situation and resources.

Enter a number to learn more about one of these points on **service diversification** or ask a new question.

Q10. *How do I increase the returns on my cafeteria?*

A10. Understanding your business needs is crucial. I'm assuming your cafeteria is in Kenya, but could you provide more details? What's the size of your cafeteria, and where is it located? What's the competition like? Any specific challenges you are facing?

Based on general information, here are some strategies to increase returns:

1. ****Improve Menu****: Review your menu to ensure it appeals to local tastes and preferences. Make sure you're offering dishes that not only taste good but also have high-profit margins. Experiment with different dishes and keep track of which ones sell best.
2. ****Cost Control****: Manage food waste by implementing inventory management systems and portion control measures. This can help reduce costs and increase profitability.
3. ****Customer Experience****: Enhance the customer experience by training staff in customer service skills, improving café ambiance or cleanliness, or even offering free Wi-Fi if feasible.
4. ****Promotions & Loyalty Programs****: Implement a loyalty program or offer special promotions during slow periods to attract more customers.
5. ****Community Engagement****: Get involved with the local community through sponsorships or events which can enhance visibility and build a loyal customer base.

Enter a number to learn more about one of these points on **increasing cafeteria returns** or ask a new question.

Q11. *I have tried to offer a new special meal menu in my restaurant and I must admit that the returns have been quite fair. Is it okay if I do a whole menu overhaul and start offering some different special meals altogether, even though it would make the price of each plate to be higher than before. I'm stuck between giving new menu or staying with current ones for the fear of losing the current clientele.*

A11. That's a great question and it's normal to feel this way when considering major changes in your business. Here are some strategies you can consider:

1. ****Test the Waters****: Before overhauling the entire menu, consider introducing new special meals one at a time. This will allow you to gauge customer response without risking too much.
2. ****Customer Feedback****: Ask your customers directly for their opinion on potential new dishes or even higher prices. You could use surveys or informal conversations.
3. ****Market Research****: Look at what similar restaurants in your area are doing successfully, especially those with a similar clientele.
4. ****Pricing Strategy****: Consider if there's room to slightly increase prices of current dishes instead of introducing more expensive ones that might deter customers.
5. ****Special Events****: Host special events where only the new menu is available, like themed nights or tasting events, to introduce customers to these dishes and gather feedback.

Enter a number to learn more about one of these points on **menu overhaul** or ask a new question.

Q12. 1

A12. Great, you're interested in the strategy of testing the waters with new special meals. Here are some additional tips on how to do that:

1. ****Pilot Dish****: Start by introducing one or two new dishes as a special offer and observe customer reactions.
2. ****Promote****: Make sure your customers know about the new dish through table tents, menu boards, or even verbally when they're ordering.
3. ****Feedback****: Ask for feedback directly from those who try the new dishes. This could be done informally by waitstaff or through a small comment card.
4. ****Evaluate****: Track sales of the new dishes and compare them to your regular items. If they perform well, consider adding them permanently.
5. ****Adjust****: Based on feedback and sales data, adjust the dish if needed before fully integrating it into your menu.

Enter a number to learn more about one of these points on **testing new special meals** or ask a new question.

Q13. *As for the customer feedback, sometimes it may not be reliable since some of the clients we host are passengers on transit so some opinion they may give may not apply to everyday's consideration.*

A13. You're correct, transient customers may provide feedback that doesn't necessarily reflect the preferences of your regular clientele. Here are some strategies to ensure you're getting reliable feedback:

1. ****Target Regular Customers****: Make a point to solicit feedback from your regulars who understand your current menu and can provide insights into potential changes.
2. ****Digital Feedback****: If you have an online presence (e.g., social media), use these platforms to gather feedback as people might be more open and detailed in their responses.
3. ****Feedback Cards****: Leave these at tables or the checkout counter for customers to fill out before they leave. This will give you a wider range of opinions.

4. **Observation**: Watch what dishes get ordered most often and which ones come back with leftovers – this can be very telling about what’s working and what isn’t.
5. **Trial Periods**: As mentioned earlier, introduce new items slowly and see how they perform over a set period before making permanent decisions.

Enter a number to learn more about one of these points on *collecting reliable customer feedback* or ask a new question.

Q14. *A shift in business strategies may sometimes backfire. Implementing new ideas may fail the business often. How can I avoid the setbacks off trying new ideas and taking risks?*

A14. Trying new ideas and taking risks is a part of business growth, but it’s important to manage these risks wisely. Here are some strategies:

1. **Market Research**: Before implementing a new idea, conduct thorough market research to understand its potential impact on your customers and competitors.
2. **Pilot Testing**: Instead of implementing changes across the entire business, try them in a small controlled environment first to gauge their effectiveness.
3. **Financial Analysis**: Consider the financial implications of the new strategy. Calculate potential returns against costs and assess if it’s financially viable.
4. **Risk Assessment**: Identify potential pitfalls or challenges with your new strategy and develop contingency plans to address them.
5. **Feedback & Review**: Regularly review the performance of the new strategy, seek feedback from customers and staff, and be ready to make adjustments as needed.

Enter a number to learn more about one of these points on *risk management when trying new business strategies* or ask a new question.

E Survey Questions and Measurement

Business performance was measured using standard survey questions that have been widely applied in developing country contexts (De Mel, McKenzie, and Woodruff, 2009). Specifically, in each of the three pre-treatment and four post-treatment performance surveys, we collected self-reported measures of weekly revenue, monthly revenue, weekly profits, and monthly profits:

- **Weekly Revenue:** Now we would like to ask you about the **total sales across all of your businesses in the last 7 days**. Please think about all the money that came into your businesses during the past 7 days, or all the money that you collected from customers BEFORE paying for any expenses and wages. How much money did you collect from all of your customers in KSH?
- **Monthly Revenue:** Now we would like to ask you about the **total sales across all of your businesses in the last 30 days**. Please think about all the money that came into your businesses during the past 30 days, or all the money that you collected from customers BEFORE paying for any expenses and wages. How much money did you collect from all of your customers in KSH?
- **Weekly Profits:** Now we would like to ask you about the **total profits** across all of your businesses in the **last 7 days**. What was the total income the business earned during the last 7 days after paying all expenses including the wages of employees, but not including any income you paid yourself. That is, what were the profits of your business in the last 7 days in KSH?
- **Monthly Profits:** Now we would like to ask you about the **total profits** across all of your businesses in the **last 30 days**. What was the total income the business earned during the last 30 days after paying all expenses including the wages of employees, but not including any income you paid yourself. That is, what were the profits of your business in the last 30 days in KSH?

In the post-treatment performance waves, we additionally state in the first performance question that “For this question and all following questions please enter ”0” if all of your businesses have closed”.

E.1 Explanation of Preferred Functional Forms

We defined four different versions of our performance index in our preanalysis plan: 95% and 99% winsorization and a log-like inverse hyperbolic sine transformation $y' = \log(y + (y^2 + 1)^{1/2})$ applied to each of these two winsorizations. In the manuscript body, we focus on the “raw” (non-log-like) performance index for two reasons.

First, recent work shows that the inverse hyperbolic sine transformation (and related log-like transformations) is highly sensitive to statistical artifacts such as the units used to measure outcomes (for example, revenue measured in individual Kenyan Shillings versus thousands of Shillings) and the presence of zeros (Chen and Roth, 2024; McKenzie, 2023). We observe this zero-sensitivity issue in our data. Figure A13 displays the distribution of our performance indices, winsorized at 95% (left) and 99% (right). The top row reports outcomes in raw units and the bottom row in log-like units. For the log-like index, zeros generate long-left-tail values below -5 standard deviations.

Second, a key focus of this study is on how AI affects performance gaps across entrepreneurs. For this purpose, changes in raw levels are more directly interpretable than proportional changes. A given percentage increase can correspond to very different absolute gains depending on whether an entrepreneur starts near the bottom or near the top of the distribution. Because we are interested in how AI shifts the level and dispersion of performance — and thus the magnitude of gaps — focusing on raw changes provides a clearer basis for comparison across the distribution.

E.2 Survey-based Mechanisms Questions and Analysis

This section provides an overview of a series of prespecified survey measures designed to explore potential mechanisms. We start by providing the verbatim survey questions for the seven sets of mechanisms:

Management practices (from McKenzie and Woodruff (2017); McKenzie, Osman, and Rahman (2023))

In the following set of questions, we will ask you about certain business practices that you may or may not be implementing in your business. Please only respond ‘Yes’ if you have

implemented this business practice at least one time in the last 3 months. It is okay if you are not doing every activity we describe and not all the practices we ask about may make sense for all businesses to do. We are just trying to assess where your business stands now.

Over the last three months...

- ...have you visited at least one of your competitor's businesses to see what prices they are charging? {No,Yes}
- ...have you visited at least one of your competitor's businesses to see what products your competitors have available for sale? {No,Yes}
- ...you asked existing customers whether there are any other products they would like your business to sell or produce? {No,Yes}
- ...have you talked with at least one former customer to find out why they have stopped buying from your business? {No,Yes}
- ...have you asked a supplier about which products are selling well in your business's industry? {No,Yes}
- ...have you ever attracted new customers with a special offer? {No,Yes}
- ...have you marketed or advertised in any form? {No,Yes}
- ...have you made use of digital social media for marketing purposes? (e.g., Facebook, WhatsApp, Twitter, LinkedIn) {No,Yes}
- ...have you attempted to negotiate with a supplier/vendor for a lower price on raw material/stock/inputs to your business? {No,Yes}
- ...did you keep track of every sale made by the business (all money into the business)? {No,Yes}
- ...did you keep track of every purchase/expenditure (all money out) made by the business? {No,Yes}
- ...did you use records regularly to know whether sales of a particular product are increasing or decreasing from one month to another? {No,Yes}
- ...did you know which goods, products, or services make the business the most profit per item selling? {No,Yes}
- ...have you worked out the cost to your business for each of your main products or services you sell? {No,Yes}
- ...did you have a written budget, which states how much is owed or should be spent each month on rent, electricity, equipment maintenance, transport, advertising, and other indirect

costs to business? {No,Yes}

- ...have you set targets for how much you will sell? {No,Yes}

Time management

- In the last seven days, how effectively did you feel you managed your time in running your business? {I managed my time running my business very poorly, Poorly, Neither poorly or well, Well, I managed my time running my business very well}
- In the last seven days, how many of your planned tasks were you able to accomplish? {I accomplished none of my planned tasks, A few of them, About half of them, Most of them, I accomplished all of my planned tasks}
- In the last seven days, how often did you find yourself working outside of your planned hours or feeling rushed? {Never worked outside planned hours or felt rushed, Rarely, Sometimes, Often, Always worked outside planned hours or felt rushed}

Innovation

- In the past 30 days, have you considered any new products, services, or business process ideas for your business? {No (I have NOT considered any new products, services, or business process ideas), Yes (I HAVE considered new products, services, or business process ideas)}
- In the past 30 days, have you tested or experimented with any new products, services, or business process ideas for your business? {No (I have NOT tested or experimented with new products, services, or business process ideas), Yes (I HAVE tested or experimented with new products, services, or business process ideas)}
- In the past 30 days, have you implemented any new products, services, or business process ideas for your business? {No (I have NOT implemented with new products, services, or business process ideas), Yes (I HAVE implemented with new products, services, or business process ideas)}

Technology use

- How frequently do you use technology (like computers, applications, software) in running your business? I never use technology (like computers, applications, software) in running my business {Rarely, Sometimes, Often, I very often use technology (like computers, applications, software) in running my business}

- To what extent do you think technology has impacted your business performance like your growth in profits or customer base? {Technology has had no effect, Technology has had a small effect, Technology has had a moderate effect, Technology has had a large effect, Technology has had a very large effect}
- In the last 30 days, have you started using any new technologies in your business? {No (no new technologies), Yes, one new technology, Yes, two new technologies, Yes, three or more new technologies}

Information: Accessibility and Quality

- In the last 30 days, how easy has it been for you to get advice, mentorship, or information for your business when you need it? {Very difficult to get advice, mentorship, or information, Difficult, Neutral, Easy, Very easy to get advice, mentorship, or information}
- In the last 30 days, how would you rate the quality of advice, mentorship or information you received for your business (across all sources you receive advice, mentorship, and information from)? {The quality of advice, mentorship, and information I receive is Very poor, Poor, Neutral, Good, The quality of advice, mentorship, and information I receive is very good}

Information: Stigma and negative judgment

- In the past 30 days, how often have you felt judged or looked down upon when asking for advice, mentorship, or information for your business? {I have always felt judged or looked down upon when asking for advice, mentorship, or information for my business, Almost always, Sometimes, Almost never, I have never felt judged or looked down upon when asking for advice, mentorship, or information for my business}
- In the past 30 days, has fear of negative judgment ever prevented you from seeking advice or mentorship or information for your business? {Fear of negative judgment has always prevented me from seeking advice or mentorship or information for my business, Almost always, Sometimes, Almost never, Fear of negative judgment has never prevented me from seeking advice or mentorship or information for my business}

Information: Relevance of mentorship/advice

- In the past 30 days, is the advice, mentorship, or information you have received directly

relevant to your current business needs? {The advice I received is never relevant to my current business needs, Rarely, Sometimes, Often, The advice I received is always relevant to my current business needs}

- In the past 30 days, how frequently have you found the advice, mentorship or information you received to be modern and up-to-date with current business trends and practices? {The advice I received is never modern and up to date, Rarely, Sometimes, Often, The advice I received is always modern and up to date}

Analysis of Survey-Based Mechanism Questions

This section provides an overview of the analysis of the families of survey outcomes outlined in the previous section:

- Management practices (sixteen items)
- Time management (three items)
- Technology use (three items)
- Innovation (three items)
- Information: Relevance (two items)
- Information: Access stigma (two items)
- Information: Accessibility and quality (two items)

For each set of outcomes we create a simple aggregate index. Specifically, for each survey question within each outcome family, we first standardized the question relative to the control group. We then constructed the index by taking the simple average across these variables in each survey wave. The analysis of these mechanism questions is presented in Figure A10. In this figure, we apply the same strategy used to identify average and heterogeneous treatment effects on business performance in the main text. Across these prespecified survey mechanisms, we observe no consistent patterns in terms of either average or heterogeneous treatment effects.

One possible explanation for the lack of treatment effects is a unique feature of the AI assistant: unlike standardized business training programs that provide identical content to all participants, the assistant's responses are tailored to the topics raised by each entrepreneur. This customization may lead to variation in mechanisms and outcomes across users that is

not captured by our prespecified survey measures. The lack of clear patterns in the survey data motivates our shift to analyzing the text of entrepreneur-AI conversations to better understand the sources of heterogeneity in the main results.

F Ethical Considerations

In our field experiment, we found no overall impact on business performance from the treatment intervention. However, we did detect heterogeneous treatment effects with respect to initial performance levels; high performers experienced a performance increase due to AI advice, while low performers experienced a performance decrease. These unexpected negative results gave us pause. Given that our treatment appears to have led to worse performance for some entrepreneurs in our sample, in this section, we outline why our experiment is ethical, above and beyond having received approval from and followed the policies of the UC Berkeley’s Institutional Review Board.

First, according to Asiedu et al. (2021), a key ethical consideration is policy equipoise, which involves meaningful uncertainty about the efficacy of each treatment arm. Indeed, if there was consensus that one treatment arm would be superior to the other, then it would *not* be ethical to randomize it unless the treatment was costly and therefore scarce. In the context of our experiment, there was uncertainty regarding participants’ benefits from each arm of the study. To the best of our knowledge, we are the first to test whether AI advice can increase entrepreneurial performance, especially in a context like Kenya. That said, recent research on generative AI and research on advising and mentoring between entrepreneurs suggested that an AI assistant had the potential to increase business performance. While we believed there were benefits, there was no consensus at the time we deployed the experiment that one treatment would be superior to the other. Given the presence of equipoise, we argue it was more ethical to test the impact of AI with a few hundred entrepreneurs in a well-designed study, rather than not testing and making policy decisions that could a) deny millions of entrepreneurs access to a useful tool if the treatment effect was positive or b) encourage adoption of a technology that hurts productivity by millions of entrepreneurs if the treatment effect was negative (Glennerster and Takavarasha, 2014).

Beyond the presence of equipoise, we also ensured all participants in our study were well-informed and free to decide whether to participate or not. Our participants were also free to decide whether to continue using our AI assistant, and it was easy to stop using it, consistent with the ethical standards around participant choice.

While our experiment revealed that one particular subgroup did not benefit from access to the AI assistant, negative treatment effects for subpopulations commonly occur in experiments conducted in the social and medical sciences (Karlan, Knight, and Udry, 2015). Given the growing use of generative AI, understanding the impact of such tools on society is especially valuable.

G Experimental Spillovers

In this appendix, we examine the plausibility of two types of spillovers that could impact our results: treatment access spillovers, where non-treated individuals might gain access to the AI assistant, and performance spillovers (Holtz et al., 2025), in which performance improvements (or declines) for one individual come at the direct cost (or benefit) of another.

Treatment Access Spillovers

The AI assistant’s WhatsApp contact number is publicly available, allowing anyone to text in. Without strict controls, this raises the possibility that untreated participants could interact with the AI assistant. To address this, we developed our AI assistant in such a way that it only responds to messages from individuals on a prespecified access list. Only participants in the AI assistant condition were added to this list.

Performance Spillovers

Another potential concern is performance spillovers within our experimental sample. Specifically, the heterogeneous treatment effects we observe might be driven by high-performing entrepreneurs capturing profits or revenue from low-performing entrepreneurs in the same sample. We consider this unlikely because our recruitment Meta advertisement targeted a

broad audience and drew participants from across Kenya, a country with over 50 million people and high Facebook penetration (Wamuyu, 2020). This recruitment strategy resulted in a geographically diverse sample. Figure A1 shows the distribution of respondents by Kenyan county, and highlights the wide variation in location.

Despite the fact that it is unlikely that spillovers would be affecting our experimental estimates, it *is* possible that performance spillovers occur between entrepreneurs in our sample and their competitors that are *not* in our sample. Keeping this in mind, the low level of technology saturation and infeasibility of the diffusion of our AI assistant are key to interpreting our results. If there were a greater level of overall generative AI saturation in Kenya, our results may not hold. For instance, it is conceivable that with high levels of generative AI adoption, firm practices could uniformly improve and/or converge; this could lead to greater price competition and lower firm profits. Conversely, widespread adoption of generative AI could drive positive performance spillovers stemming from the diffusion of best-use practices among businesses. While our experiment is unable to shed light on these dynamics, we consider this a fruitful area for future research.

H Text Analysis

This section describes our text data and analysis in detail. Subsection H.1 provides an overview of our three sources of text data (entrepreneur questions, AI responses, and descriptions of businesses changes) along with summary statistics on these datasets. Subsection H.3 describes how we use GPT-4 to develop taxonomies for classifying entrepreneur questions into 10 question topics and AI responses into 12 topics, showing that low and high performers ask about similar topics and receive similar types of advice. Subsection H.4 describes our word- and embedding element-level analyses testing whether the content of the AI’s advice causally impacts the business changes made by entrepreneurs, including a bag-of-words validation of our embedding approach. Subsection H.5 details how we constructed treatment effect embedding vectors to systematically identify representative examples of how treated versus control entrepreneurs make different business changes. Finally, Subsection H.6 describes our quantitative validation tests using keyword matching and word frequency analysis

to test the qualitative patterns identified through the treatment effect embeddings.

H.1 Overview and Summary Statistics

Our text analyses use three sets of data. First, we use data on the quantity and content of messages each entrepreneur sent to the AI assistant. These messages include any questions asked of the AI, numeric followups (e.g. the number “3” in Figure 1), and any other messages sent to the AI. Panel A of Table A5 presents basic summary statistics from these messages. For our text analyses and when estimating word embeddings, we combine the messages sent by each entrepreneur to the AI into a single text blob per entrepreneur, with summary statistics for these combined text blobs shown in Panel A of Table A11. Given that many questions are uninterpretable fragments (e.g. “3”), when classifying entrepreneurs’ questions into topics, we had two human coders read through the entire set of 4,810 messages, filtering out messages that were numeric (e.g. following up with “1”), fragments and follow-ups (e.g. “tell me more,” “Can you explain how I might do ‘3?’”), and messages not focused on business topics (e.g., a business owner who asked about how to quit smoking). The topic analysis, when applied to both questions and answers, focuses on the 1,450 conversational pairs with questions that were classified as business-relevant and therefore suitable for topic classification.

Second, we use data on the AI’s responses to entrepreneurs. Panel B of Table A5 provides summary statistics on these responses, which were often structured to provide 3-5 suggestions, as shown in Figure 1. As with the entrepreneurs’ messages, we concatenate all AI responses sent to each entrepreneur into a single text blob per entrepreneur (Panel B of Table A11) and use this combined text blob to estimate word embeddings for AI advice.

Third, we use data on business changes that each entrepreneur describes making. As part of our post-treatment surveys, we asked entrepreneurs to describe changes they made to their business during the final three survey waves, as well as any other new business activities in the last two waves. Entrepreneurs who did not make changes were not required to provide descriptions. Table 2 shows that neither treatment nor performance status predicts whether an entrepreneur describes making changes. The specific survey questions were:

- Please describe in detail any new products that you have implemented in the last 30

days. If you have not implemented any new products, you can leave this question blank.

- Please describe in detail any new services that you have implemented in the last 30 days. If you have not implemented any new services, you can leave this question blank.
- Please describe in detail any new business processes or strategies that you have implemented in the last 30 days. If you have not implemented any new business processes or strategies, you can leave this question blank.
- In the last 30 days, which change that you made to your business has had the largest impact on performance (for example, sales and profits)? Please describe it in as much detail as possible.
- In 1-2 sentences, please write the main activity of any new businesses.

As with the messages sent and advice received, we analyze the business changes made by concatenating all described changes for each entrepreneur into a single text blob and use this combined text blob to estimate word embeddings for business changes. Summary statistics for these combined business changes text blobs are shown in Panel C of Table A11.

H.2 Random Forest Classification Details

To test whether text embeddings can predict group membership (e.g., pre-treatment performance status or treatment assignment), we use random forest classifiers (Breiman, 2001). For each model, we perform 100 Monte Carlo iterations: in each iteration, we split the data 60/40 into training and test sets, a random forest is fit on the training set, and accuracy is evaluated on the held-out test set. We fit models with the `caret` package in R (Kuhn, 2008), using the default hyperparameters of 500 trees and minimum node size of 1. The optimal number of randomly sampled predictors (`mtry`) is selected via 5-fold cross-validation within the training set, evaluating values of 2, $\sqrt{p}/2$, and \sqrt{p} (with $p = 3,072$). We report the mean accuracy and the 95% empirical interval (2.5th and 97.5th percentiles) across the 100 iterations, benchmarked against the no-information rate, defined as the accuracy achieved by always predicting the majority class. A classifier must exceed the NIR to demonstrate predictive signal beyond the base rate.

H.3 Question and Answer Topic Analysis

H.3.1 Taxonomy Generation for Business Messages

To understand the topics in the messages sent to and answers generated by the AI, we follow an approach outlined in Choi et al. (2023) and use GPT-4 to identify topics within the business-related questions posed by entrepreneurs to the AI assistant. This process has four stages:

Stage 1: As described above, we had two human coders read through the entire set of messages, identifying questions that were business-focused, and which had sufficient context for topic classification. We focus all topic analyses on the set of 1,450 exchanges in which the question was suitable for this classification exercise.

Stage 2: Using this sample of business messages, we instructed GPT-4 to generate topics to classify the messages in the corpus using the prompt below.

Listed below are queries submitted by SME entrepreneurs in Kenya to an AI Business Mentor designed to help them in their daily business operations.

The AI Business Mentor is designed to provide guidance and support to SME entrepreneurs in Kenya by answering their queries and providing them with valuable insights and advice. For each query presented, provide [**num.topics**] high-level business topics in less than 4 words each.

For example, if the query is “How can I improve my sales?”, the corresponding 4 topics could be [“Sales Strategy”, “Customer Acquisition”, “Marketing Campaigns”]

Return the topics for each query in a JSON format with the queryID, query and the corresponding topics.

The query list can be found below:

```
### query_list ###
```

We ran this process multiple times, beginning with 3 topics per query and incrementally increasing to 20 topics per query. We then combined all topics generated across these runs to form a comprehensive set of 207 distinct topics. For robustness, we randomly selected approximately 10% of the topic assignments to messages and manually reviewed them to verify the topics were sensible.

Stage 3: We then took the set of 207 topics and instructed GPT-4 to merge or eliminate redundant topics, narrowing them down to a more concise list. We manually refined this list to remove overly specific topics, settling on a set of 10 topics:

1. Financial Management
2. Marketing Strategies

3. Risk Management
4. Customer Acquisition & Retention
5. Business Planning & Growth
6. Technology Integration & Digital Marketing
7. Inventory & Operations Management
8. Market Research & Analysis
9. Livestock & Farm Management
10. Entrepreneurial Opportunities & Business Ideas

Stage 4: Using these ten high-level topics, we instructed GPT-4 to classify each business message into one of these categories. Messages could only be classified into a single topic. We used the following prompt to classify these messages:

Listed below are queries submitted by SME entrepreneurs in Kenya to an AI Business Mentor designed to help them in their daily business operations. The AI Business Mentor is designed to provide guidance and support to SME entrepreneurs in Kenya by answering their queries and providing them with valuable insights and advice. Provide a classification of each query using ONE high-level business topic from the list of topics below. For example, if the query is “How can I improve my sales?”, the corresponding topic could be “Marketing Strategies” or “Customer Acquisition & Retention.”

Here is the list of topics to choose from:

- Financial Management
- Marketing Strategies
- Risk Management
- Customer Acquisition & Retention
- Business Planning & Growth
- Technology Integration & Digital Marketing
- Inventory & Operations Management
- Market Research & Analysis
- Livestock & Farm Management
- Entrepreneurial Opportunities & Business Ideas

Return the topic for each query in a JSON format with the keys: `queryID`, `query`, and `topic`.

You are ONLY permitted to respond with one of these topics and must NEVER return any other text outside this list.

(query)

We then used the topics generated from this procedure to create Figure A14 which shows the distribution of topics by performance status. The topic distributions are very similar for high- and low-performers.

H.3.2 Taxonomy Generation for AI Mentor Responses

Following the process outlined above, we also generated a set of topics for the messages generated by the AI assistant in response to the entrepreneurs’ business messages. The result of the first three steps is a list of 12 topics:

1. Financial Management
2. Marketing Strategies
3. Business Growth & Expansion

4. Operational Efficiency
5. Agriculture
6. Customer Engagement
7. Starting a Business
8. Product Differentiation
9. Market Research & Analysis
10. Risk Management
11. E-commerce Platforms
12. Human Resources

We then used the following system prompt to classify each response from the AI assistant into a single topic:

Listed below are responses from an AI Business Mentor designed to help SME entrepreneurs in Kenya with their daily business operations. Classify each response using ONE high-level business topic from the list below. For example, if the response discusses strategies for managing cash flow or bookkeeping, the topic would be “Financial Management.” If the response advises on how to attract and retain customers, the topic would be “Customer Engagement.”

Here is the list of topics to choose from:

- Financial Management
- Marketing Strategies
- Business Growth & Expansion
- Operational Efficiency
- Agriculture
- Customer Engagement
- Starting a Business
- Product Differentiation
- Market Research & Analysis
- Risk Management
- E-commerce Platforms
- Human Resources

Return the topic for the response in a JSON format with the keys: `queryID`, `response`, and `topic`. You are ONLY permitted to respond with one of these topics and must NEVER return any other text outside this list. (response)

Figure A15 shows the distribution of response topics for low and high performers. The distribution largely mirrors what we see in Figure A14, in that we again find similar response patterns across the groups.

H.4 Bag of Words Analysis

As a robustness check for our embedding-based approach for measuring whether entrepreneurs follow AI advice, we conducted a complementary “bag of words” style analysis. Table A6 shows that when one of the more than 12,000 non-stop words in the AI’s responses is present in an entrepreneur’s advice text, it is associated with a higher probability that the same word appears in the entrepreneur’s business changes text by 2.80 percentage points ($p < 0.001$, 95% CI = [2.32, 3.29]), while controlling for the question text. However, this approach

rests on the strong assumption that words impact only themselves (e.g., “price” in advice affects only “price” in business changes, not related terms like “discount”) (Grimmer and Stewart, 2013). This “bag of words” assumption is clearly violated when, for example, advice about “pricing” leads entrepreneurs to describe “discounts.” Because of this, we prefer our embedding-based approach, which addresses these concerns by capturing semantic relationships across the full 3,072-dimensional vector space, accounting for edge cases where entrepreneurs use semantically similar but different words.

H.5 Treatment Effect Embedding Methodology

As mentioned in the main text, we developed a novel empirical strategy to identify the types of text changes that best represent the average differences between groups. This strategy involves first calculating the average difference in text embeddings between treatment and control groups for low performers and then for high performers. By comparing these embedding vectors, we capture how the AI-generated advice systematically impacted the types of changes low-performing entrepreneurs report making to their business. Specifically, we calculate the average business change embedding vector for low-performing treated entrepreneurs, calculate the average business change embedding vector for low-performing control entrepreneurs, and then compute the element-wise difference between treated and control averages (the “treatment effect embedding”). We then do exactly the same procedure for high performers.

Because these 3,072-dimension vectors are not easily interpretable on their own, we then take pairs of control and treated business changes text, calculate the difference in the embedding vector, and identify individual text pairs that are most similar (using cosine similarity distance) to the average difference between groups. In Table A9 and Table A10 we qualitatively illustrate how treated versus control entrepreneurs differ in the types of business changes they make. Both tables show example pairs of treated versus control text sampled from the top 1% of most cosine similar text pairs to the high- or low-performer treatment effect embedding.

H.6 Implementation of Word Matching Analyses

For our word matching analysis we construct three measures. First, to test whether the entrepreneur mentions working with and learning from the AI, we take the raw business changes text and match it with the following regular expression pattern:

```
AI|[Cc]hat|[Gg][Pp][Tt]|[Aa][Ii]  
[Mm]entor|[Aa]rtificial [Ii]ntelligence
```

We classify an entrepreneur who matches this pattern as having described working and learning from the AI. Second, to test if the entrepreneur references the PDF guide we distributed to the control group, we test if the raw business changes text matches the following regular expression pattern:

```
[Gg]uide|[Pp][Dd][Ff]|[Mm]anual
```

Third, to test if the entrepreneur offered discounts or invested in advertising we match whether the raw business changes text matches the following regular expression:

```
[Ff]ree|[Dd]iscount|[Rr]educ|[Oo]ffer|[Aa]dver
```

We deliberately do not match against “marketing” as many entrepreneurs describe how they tried non-paid marketing. We include “free” because many entrepreneurs mention offering free delivery and other services that implicitly serve as a discount.¹²

We estimate linear probability models regressing the binary indicator for whether an entrepreneur’s text matched each regular expression on treatment assignment and baseline performance status. For these measures, we assume that entrepreneurs who do not describe making any business changes do not match any of the regular expressions, since activities such as working with the AI, building on the guide, or implementing pricing strategies necessarily require making business changes. Table A7 presents regression estimates for AI mentions (Panel A), ILO business guide references (Panel B), and discounting/advertising strategies (Panel C). In Panel A, we find that those in the AI assistant condition report

¹²While these regular expressions provide a systematic approach to quantifying the patterns observed in Table A9 and Table A10, they necessarily involve some measurement error inherent to keyword-based classification methods.

learning from AI, while we find no such evidence of engagement with the PDF guides in Panel B, suggesting that the guides served as a placebo control. Finally, Panel C presents evidence that the AI assistant leads initially low-performing entrepreneurs to be more likely to engage in generic strategies like discounts and advertising.

The final measure counts the number of uncommon words used in an entrepreneur’s business changes text. To create this measure, we removed capitalization, tokenized words, and removed stop words from the business changes text. We then classified words as uncommon if they appear in 5 or fewer business change descriptions. These words include “150cc,” “arcade,” “biodegradable,” “detergent,” “headphones,” “peanut,” “rabbit” and “uganda.” Common words include “advertising,” “customers,” “market,” “price,” and “vegetables.”

Table A8 presents Poisson regressions testing whether treatment and performance status predict the count of uncommon words used. We use Poisson regressions because the outcome is a highly skewed count variable. Since entrepreneurs who write longer responses will mechanically use more uncommon words, we control for the total word count in the business changes text. This analysis is restricted to entrepreneurs who provided at least one analyzable word in their business changes text.

References

- Asiedu, Edward, Dean Karlan, Monica Lambon-Quayefio, and Christopher Udry. 2021. “A call for structured ethics appendices in social science papers.” *Proceedings of the National Academy of Sciences* 118 (29):e2024570118.
- Benjamini, Yoav and Yosef Hochberg. 1995. “Controlling the false discovery rate: a practical and powerful approach to multiple testing.” *Journal of the Royal statistical society: series B (Methodological)* 57 (1):289–300.
- Breiman, Leo. 2001. “Random forests.” *Machine learning* 45 (1):5–32.
- Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2025. “Generative AI at work.” *The Quarterly Journal of Economics* 140 (2):889–942.
- Chen, Jiafeng and Jonathan Roth. 2024. “Logs with zeros? Some problems and solutions.” *The Quarterly Journal of Economics* 139 (2):891–936.
- Choi, Jun Ho, Oliver Garrod, Paul Atherton, Andrew Joyce-Gibbons, and Daniel Mason-Sesay, Miriam & Björkegren. 2023. “Are LLMs Useful in the Poorest Schools? TheTeacherAI in Sierra Leone.” *Preprint, submitted. October 4th, 2023, <https://arxiv.org/abs/2310.02982>*.
- De Mel, Suresh, David J McKenzie, and Christopher Woodruff. 2009. “Measuring microenterprise profits: Must we ask how the sausage is made?” *Journal of development Economics* 88 (1):19–31.
- Glennerster, Rachel and Kudzai Takavarasha. 2014. *Running randomized evaluations: A practical guide*. Princeton University Press.
- Grimmer, Justin and Brandon M Stewart. 2013. “Text as data: The promise and pitfalls of automatic content analysis methods for political texts.” *Political analysis* 21 (3):267–297.
- Holtz, David, Felipe Lobel, Ruben Lobel, Inessa Liskovich, and Sinan Aral. 2025. “Reducing interference bias in online marketplace experiments using cluster randomization: Evidence from a pricing meta-experiment on airbnb.” *Management Science* 71 (1):390–406.

- Kamau, Paul, Laura Barasa, Gedion Onyango, and Oscar Otele. 2024. "Summary of results: Afrobarometer Round 10 survey in Kenya, 2024." Survey summary, Institute for Development Studies, University of Nairobi & Afrobarometer. <https://www.afrobarometer.org/wp-content/uploads/2024/10/Summary-of-results-Kenya-Afrobarometer-R10-template-Eng-16oct24-for-upload.pdf>.
- Karlan, Dean, Ryan Knight, and Christopher Udry. 2015. "Consulting and capital experiments with microenterprise tailors in Ghana." *Journal of Economic Behavior & Organization* 118:281–302.
- Kuhn, Max. 2008. "Building predictive models in R using the caret package." *Journal of statistical software* 28 (5):1–26.
- McKenzie, David. 2023. "Interpreting treatment effects on an inverse hyperbolic sine outcome variable and alternatives." *Development Impact* Accessed: 2023-11-29.
- McKenzie, David, Adam Osman, and Aminur Rahman. 2023. "Training and subsidies vs pay-for-results in spurring digital marketing take-up and small firm growth." *Journal of Development Economics Registered Report Proposal*.
- McKenzie, David and Christopher Woodruff. 2017. "Business practices in small firms in developing countries." *Management Science* 63 (9):2967–2981.
- Mehmood, Muhammad Zia. 2024. "Short Messages Fall Short for Micro-Entrepreneurs: Experimental Evidence from Kenya." *Preprint September, 2024*, <https://www.theigc.org/sites/default/files/2025-03/Mehmood-Working-Paper-September-2024.pdf>.
- Wamuyu, PK. 2020. "The Kenyan Social Media Landscape: Trends and Emerging Narratives, 2020." SIMElab, Nairobi, <https://www.scribd.com/document/609901119/Kenya-Social-Media-Landscape-Report-2020>.