

PHD_THESIS_Adelaide_Fabbi.pdf

 Institut Polytechnique de Paris

Détails du document

ID de la copie

trn:oid::2945:408092279

Date de la copie

17 août 2026, 12:00 UTC+2

Date de téléchargement

17 août 2026, 12:44 UTC+2

Nom du fichier

PHD_THESIS_Adelaide_Fabbi.pdf

Taille du fichier

7.0 MB

pages

mots

caractères

12% Similarité globale

Total combiné de toutes les correspondances, y compris les sources qui se recoupent, po...

Groupes de correspondances

- 565 Pas de citations 10%**
Correspondances sans citation dans le texte ni guillemets
- 36 Guillemets manquants 1%**
Correspondances restant très similaires au support source
- 29 Citation manquante 1%**
Correspondances avec guillemets, mais pas de citation dans le texte
- 0 Citations 0%**
Correspondances avec citation dans le texte, mais sans guillemets

Sources principales

- 10% Sources Internet
- 9% Publications
- 2% Travaux envoyés (copies des étudiants)

Indicateurs d'intégrité

2 indicateurs d'intégrité à vérifier

- Caractères remplacés**
299 caractères suspects sur 57 pages
Les lettres sont échangées par des caractères similaires provenant d'un au
- Texte masqué**
40 caractères suspects sur 2 pages
Le texte est modifié pour se fondre dans le fond blanc du document.

Les algorithmes de notre système examinent un document en profondeur pour détecter toute incohérence qui le distinguerait d'une copie normale. Si nous remarquons quelque chose d'étrange, nous le signalons pour que vous puissiez l'examiner.

Un indicateur n'est pas nécessairement le signe d'un problème. Cependant, il est recommandé d'effectuer un examen plus approfondi.

Groupes de correspondances

- 565** Pas de citations 10%
Correspondances sans citation dans le texte ni guillemets
- 36** Guillemets manquants 1%
Correspondances restant très similaires au support source
- 29** Citation manquante 1%
Correspondances avec guillemets, mais pas de citation dans le texte
- 0** Citations 0%
Correspondances avec citation dans le texte, mais sans guillemets

Sources principales

- 10% Sources Internet
- 9% Publications
- 2% Travaux envoyés (copies des étudiants)

Sources principales

Sources ayant le nombre le plus élevé de correspondances dans la copie. Les sources qui se recoupent ne sont pas affic...

1	Internet	www.econstor.eu	<1%
2	Internet	www.rfberlin.com	<1%
3	Internet	conference.iza.org	<1%
4	Internet	theses.hal.science	<1%
5	Internet	eprints.lse.ac.uk	<1%
6	Internet	tesionline.unicatt.it	<1%
7	Internet	research-information.bris.ac.uk	<1%
8	Internet	www.sciencespo.fr	<1%
9	Internet	real.mtak.hu	<1%
10	Internet	arxiv.org	<1%

11	Internet	d-nb.info	<1%
12	Publication	Jonas Cederlöf, Martin Söderström, Johan Vikström. "The Role of Caseworkers: Jo...	<1%
13	Internet	docs.iza.org	<1%
14	Internet	violasalvestrini.weebly.com	<1%
15	Internet	ifs.org.uk	<1%
16	Internet	eprints.whiterose.ac.uk	<1%
17	Internet	en.rwi-essen.de	<1%
18	Internet	benny.aeaweb.org	<1%
19	Internet	parite.math.cnrs.fr	<1%
20	Publication	Pierre Deschamps. "Gender Quotas in Hiring Committees: A Boon or a Bane for W...	<1%
21	Publication	Irena Pietrzyk, Marita Jacob. "Die Bedeutung der Treatmentprävalenz: Zur Überw...	<1%
22	Internet	www.parisschoolofeconomics.eu	<1%
23	Internet	wrap.warwick.ac.uk	<1%
24	Internet	www.cesifo.de	<1%

25	Internet	cris.maastrichtuniversity.nl	<1%
26	Internet	stax.strath.ac.uk	<1%
27	Internet	dspace.library.uu.nl	<1%
28	Publication	Ellora Derenoncourt, David Weil. "Voluntary Minimum Wages: The Local Labor M...	<1%
29	Publication	Lobel, Felipe. "Essays on Public Economics", University of California, Berkeley, 2024	<1%
30	Internet	escholarship.org	<1%
31	Internet	www.nber.org	<1%
32	Internet	publishup.uni-potsdam.de	<1%
33	Internet	scholarworks.umass.edu	<1%
34	Internet	www.iwh-halle.de	<1%
35	Internet	ideas.repec.org	<1%
36	Internet	www.hhs.se	<1%
37	Internet	eprints.gla.ac.uk	<1%
38	Publication	Amelie Schiprowski. "The Role of Caseworkers in Unemployment Insurance: Evide...	<1%

39	Internet	financedocbox.com	<1%
40	Internet	dash.harvard.edu	<1%
41	Internet	iris.unito.it	<1%
42	Internet	charly.wsr.ac.at	<1%
43	Internet	www.csef.it	<1%
44	Publication	Powell, Tyler. "Three Essays on the Economics of Education.", Michigan State Univ...	<1%
45	Internet	madoc.bib.uni-mannheim.de	<1%
46	Publication	José J. Domínguez. "Diversified committees in hiring processes: Lab evidence on g...	<1%
47	Publication	Li, Xuan. "Three Essays in Applied Microeconomics.", Boston University	<1%
48	Publication	Sondergeld, Virginia. "Women in Management: Essays in Gender and Labor Econo...	<1%
49	Internet	city-test.eprints-hosting.org	<1%
50	Internet	www.cesifo.org	<1%
51	Publication	Blau, Francine D.. "The Economics of Women, Men, and Work", Oxford University ...	<1%
52	Internet	www.pure.ed.ac.uk	<1%

53	Publication	"Neoplasms", Wiley, 2011	<1%
54	Internet	www.cambridgescholars.com	<1%
55	Internet	www.jstor.org	<1%
56	Publication	Qiu, Christina M.. "Essays on Technology and Spatial Economics.", Yale University	<1%
57	Internet	dergipark.org.tr	<1%
58	Internet	papers.saraspaziani.com	<1%
59	Publication	Rodriguez-Martinez, Pedro. "Essays in the Economics of Crime.", The George Was...	<1%
60	Internet	econpapers.repec.org	<1%
61	Internet	export.arxiv.org	<1%
62	Internet	haanwinckel.github.io	<1%
63	Internet	www.datocms-assets.com	<1%
64	Internet	dokumen.pub	<1%
65	Internet	maartengoos.com	<1%
66	Internet	pure.au.dk	<1%

67	Publication	Alan Manning, Barbara Petrongolo. "How Local Are Labor Markets? Evidence fro...	<1%
68	Publication	De Paola, Maria, and Vincenzo Scoppa. "Gender Discrimination and Evaluators' Ge...	<1%
69	Internet	epub.uni-bayreuth.de	<1%
70	Internet	files.eric.ed.gov	<1%
71	Internet	lucris.lub.lu.se	<1%
72	Publication	Cui, Jingyi. "Information and Matching in Digital Labor Markets.", Yale University	<1%
73	Publication	Michael C. Knaus, Michael Lechner, Anthony Strittmatter. "Heterogeneous Emplo...	<1%
74	Internet	econweb.umd.edu	<1%
75	Internet	jhr.uwpress.org	<1%
76	Internet	maritrehavi.github.io	<1%
77	Internet	www.duo.uio.no	<1%
78	Internet	www.snf.ch	<1%
79	Publication	Maria De Paola, Michela Ponzio, Vincenzo Scoppa. "Are Men Given Priority for Top ...	<1%
80	Publication	Thomas Le Barbanchon, Roland Rathelot, Alexandra Roulet. "Gender Differences i...	<1%

81	Internet	iris.unica.it	<1%
82	Publication	René Böheim, Rainer Eppel, Helmut Mahringer. "The impact of lower caseloads in...	<1%
83	Publication	Viola Corradini, Lorenzo Lagos, Garima Sharma. "Collective Bargaining for Wome...	<1%
84	Publication	Amelie Schiprowski. "The Role of Caseworkers in Unemployment Insurance: Evide...	<1%
85	Publication	Srivastava, Kartik. "Essays on Labor Market Frictions in Developing Countries.", H...	<1%
86	Publication	Struck, Alexander. "Dependencies in Topic Delineation and Tracking", Humboldt ...	<1%
87	Publication	THOMAS W. BAGSHAWE. "Notes on the Habits of the Gentoo and Ringed or Antar...	<1%
88	Internet	scholar.harvard.edu	<1%
89	Internet	www.stephenconnell.org	<1%
90	Publication	Ana M. Silva, José M. Arranz, Carlos García-Serrano. "Evaluating the impact of the ...	<1%
91	Publication	Blanca Casares-Guillén, Serafín Pazos-Vidal. "The role of managing authorities in ...	<1%
92	Internet	business.purdue.edu	<1%
93	Internet	gwern.net	<1%
94	Internet	www.granted-project.eu	<1%

95	Publication	Byambasuren, Tsenguunjav. "Essays on Labor Market Power.", Cornell University	<1%
96	Publication	Roach, M. A.. "Indigent Defense Counsel, Attorney Quality, and Defendant Outco...	<1%
97	Internet	gender-ict.net	<1%
98	Internet	mail.e-jlia.com	<1%
99	Internet	www.conftool.pro	<1%
100	Publication	Kazuki Nakajima, Takayuki Mizuno. "Researcher population pyramids: Tracking d...	<1%
101	Publication	Pia Homrighausen, Michael Oberfichtner. "Do caseworker meetings prevent une...	<1%
102	Publication	Trine Filges, Geir Smedslund, Anne-Marie Klint Jørgensen. "Active Labour Market ...	<1%
103	Internet	ericwilken.com	<1%
104	Internet	spaces-cdn.owlstown.com	<1%
105	Publication	Affonso Peyre, Agustina. "Essays in Economics of Education.", Arizona State Unive...	<1%
106	Publication	Thomas Le Barbanchon, Julien Sauvagnat. "Electoral Competition, Voter Bias, and...	<1%
107	Internet	digitalassets.lib.berkeley.edu	<1%
108	Internet	ftp.iza.org	<1%

109	Internet	haochenzhang.weebly.com	<1%
110	Travaux envoyés	ip-paris on 2024-09-27	<1%
111	Internet	serval.unil.ch	<1%
112	Internet	waynesandholtz.com	<1%
113	Publication	Caique Melo, Robson Tigre. "Are Educated Candidates Less Corrupt Bureaucrats? ...	<1%
114	Publication	Carlos Carrillo-Tudela, Hermann Gartner, Leo Kaas. "Recruitment Policies, Job-Filli...	<1%
115	Publication	Caspar Rose. "Chapter 52 Good Faith and Game Theory", Springer Science and Bu...	<1%
116	Publication	Conny Wunsch. "Optimal Use of Labor Market Policies: The Role of Job Search Ass...	<1%
117	Publication	Sanghee Park. "Seeking changes in ivory towers: The impact of gender quotas on ...	<1%
118	Publication	Chaturvedi, Sugat. "Representation in Politics and Labor Markets", Indian Statisti...	<1%
119	Publication	McQuillan, Casey C.. "The Causes and Consequences of Incomplete Unemployme...	<1%
120	Publication	Mehrotra, Anay. "Learning Theory in the Wild: Foundations of Missing Data and L...	<1%
121	Publication	Michael Bradley, Dong Chen. "Does Board Independence Reduce the Cost of Debt...	<1%
122	Publication	Michael C. Knaus, Michael Lechner, Anthony Strittmatter. "Heterogeneous Emplo...	<1%

123	Publication	Tomforde, Christine Emily. "Class Frequency and Timing: A Comparison of Acade...	<1%
124	Internet	deepai.org	<1%
125	Internet	huemaurice5.blogspot.com	<1%
126	Internet	laurahospido.com	<1%
127	Publication	Ciasullo, Ludovica. "Essays in Labor Economics", New York University, 2025	<1%
128	Publication	Rahul Hemrajani, Tony Hobert. "The Effects of Decision Fatigue on Judicial Behavi...	<1%
129	Internet	discovery.ucl.ac.uk	<1%
130	Internet	onlinelibrary.wiley.com	<1%
131	Internet	tel.archives-ouvertes.fr	<1%
132	Internet	www.ifau.se	<1%
133	Internet	www.researchgate.net	<1%
134	Publication	B. McCall, J. Smith, C. Wunsch. "Government-Sponsored Vocational Education for ...	<1%
135	Publication	Brandon Lehr. "Behavioral Economics - Evidence, Theory, and Welfare", Routledg...	<1%
136	Publication	Celeste K. Carruthers, Shaun Dougherty, Thomas Goldring, Daniel Kreisman, Rod...	<1%

137	Publication	Khamitov, Koblandy. "Interdisziplinäre Behandlung der Malignen Thoraxwandtu...	<1%
138	Publication	McVicar, D.. "Job search monitoring intensity, unemployment exit and job entry: ...	<1%
139	Publication	Timothy Besley, Olle Folke, Torsten Persson, Johanna Rickne. "Gender Quotas and...	<1%
140	Internet	joshuatdean.com	<1%
141	Internet	www.christinevalente.com	<1%
142	Internet	www.iza.org	<1%
143	Internet	www.science.org	<1%
144	Internet	www.sciendo.com	<1%
145	Internet	www.repo.uni-hannover.de	<1%
146	Publication	Sethi, Ritika. "Equity and Efficiency in Education", Rice University	<1%
147	Publication	Tuomas Eerola, Anna M. Czepiel. "The Gendered Nature of Authorship in Music Ps...	<1%
148	Internet	doczz.net	<1%
149	Internet	www.mdc.e.u-tokyo.ac.jp	<1%
150	Publication	Chung, Su Hwan. "Three Essays in Labor Economics", Michigan State University, 2...	<1%

151	Publication	Junming Huang, Alexander J. Gates, Roberta Sinatra, Albert-László Barabási. "Hist...	<1%
152	Publication	Rongkang Pei, Zeyu Lyu, Guolong Wang, Zhichao Wang, Maoxin Ye, Xiaoguang Fa...	<1%
153	Internet	codytuttle.github.io	<1%
154	Internet	gupea.ub.gu.se	<1%
155	Internet	www.hbs.edu	<1%
156	Publication	Beata Woźniak-Jęchorek, Kamilla Marchewka-Bartkowiak. "Digital Labour Market...	<1%
157	Publication	Chang He, Chao Kevin Li, Gary S. Monroe, Yi Si. "Diversity of Signing Auditors and ...	<1%
158	Publication	Thornton, Annabel. "Behavioural Responses to Imperfect Information: Evidence F...	<1%
159	Internet	edizioni.sns.it	<1%
160	Internet	edoc.hu-berlin.de	<1%
161	Internet	ursinaschaede.github.io	<1%
162	Internet	visir.maxwell.vrac.puc-rio.br	<1%
163	Internet	www.povertyactionlab.org	<1%
164	Publication	"Economics of Education", Springer Science and Business Media LLC, 2025	<1%

165	Publication	"International Handbook of Universities 2019", Springer Science and Business Me...	<1%
166	Publication	Marco Nieddu, Lorenzo Pandolfi. "The effectiveness of promotion incentives for p...	<1%
167	Internet	boris.unibe.ch	<1%
168	Internet	edoc.ub.uni-muenchen.de	<1%
169	Internet	link.springer.com	<1%
170	Internet	mitp.silverchair-cdn.com	<1%
171	Internet	univ-pantheon-assas.hal.science	<1%
172	Internet	www.econ.jku.at	<1%
173	Internet	www.scilit.net	<1%
174	Publication	Girum Abebe, A. Stefano Caria, Esteban Ortiz-Ospina. "The Selection of Talent: Ex...	<1%
175	Publication	Ioana Elena Marinescu, Ronald Wolthoff. "Opening the Black Box of the Matching...	<1%
176	Publication	Jens Hainmueller, Barbara Hofmann, Gerhard Krug, Katja Wolf. "Do Lower Caselo...	<1%
177	Publication	Oh, Ryungha. "Essays on Spatial Sorting and Labor Markets", Yale University, 2024	<1%
178	Publication	Radhakrishnan Gopalan, Barton Hamilton, Ankit Kalda, David Sovich. "State Mini...	<1%

179	Publication	Rafael Lalive, Jan C. van Ours, Josef Zweimüller. "The Effect of Benefit Sanctions o...	<1%
180	Publication	Zhanhan Yu, Alfonso Flores-Lagunes. "Monopsony in academia and the gender p...	<1%
181	Internet	academic.oup.com	<1%
182	Internet	events.bse.eu	<1%
183	Internet	hull-repository.worktribe.com	<1%
184	Internet	iris.unitn.it	<1%
185	Internet	italianpoliticalscience.com	<1%
186	Internet	journals.pnu.if.ua	<1%
187	Internet	jyx.jyu.fi	<1%
188	Internet	uu.diva-portal.org	<1%
189	Publication	Fabiana Roberto, Andrea Rey, Roberto Maglio, Francesco Agliata. "The academic "...	<1%
190	Publication	Grübler, Anne Sophie. "Does Gender Diversity on Boards Relate to ESG Performan...	<1%
191	Publication	Josefa Aguirre, Juan Matta, Ana María Montoya. "College peer effects, gender, an...	<1%
192	Publication	Khan, Adnan Q., Asim I. Khwaja, and Benjamin A. Olken. "Tax Farming Redux: Exp...	<1%

193	Publication	Lee, James. "Essays on the Economics of Higher Education: Transfer Students, Pee...	<1%
194	Publication	Mustafi, Priyoma. "Essays in Economics of Gender Within and Outside the Househ...	<1%
195	Publication	Roberta Ruggieri, Fabrizio Pecoraro, Daniela Luzi. "An intersectional approach to ...	<1%
196	Publication	Tandon, Shreya. "Essays on Labor Markets and Public Service Provision.", Harvard...	<1%
197	Internet	basmachelsen.nl	<1%
198	Travaux envoyés	ip-paris on 2024-10-07	<1%
199	Internet	population-economics.committee.socialpolitik.de	<1%
200	Internet	roses.univ-paris1.fr	<1%
201	Internet	trasparenza.iusspavia.it	<1%
202	Internet	ueaeprints.uea.ac.uk	<1%
203	Internet	uploads.strikinglycdn.com	<1%
204	Internet	www.bristol.ac.uk	<1%
205	Publication	Agnieszka Tycinska, Jaroslaw Drozd, Anna Fijalowska, Renata Glowczynska et al. ...	<1%
206	Publication	Anne Boschini, Anna Sjögren. "Is Team Formation Gender Neutral? Evidence from...	<1%

207	Publication	Bhuradia, Ashutosh. "Essays on Experimental Evidence for Improving Higher Edu...	<1%
208	Publication	Carlin, Wendy, Soskice, David. "Macroeconomics", Macroeconomics, 2024	<1%
209	Publication	Chattopadhyay, Shrijata. "Empirical Essays on Price Discovery Through Venture C...	<1%
210	Publication	Conny Wunsch, Felix Rochlitz, Patrick Arni. "Jobseekers' skills and job search beha...	<1%
211	Publication	Dag W. Aksnes, Shulamit Kahn, Rune Borgan Reiling, Marte E. S. Ulvestad. "Longit...	<1%
212	Publication	FRANCINE D. BLAU. "New Evidence on Gender Differences in Promotion Rates: An...	<1%
213	Publication	Isabel Fernandez-Mateo, Roberto M. Fernandez. "Bending the Pipeline? Executive...	<1%
214	Publication	Jain, Manisha. "Essays in Public Economics", The University of Wisconsin - Madison	<1%
215	Publication	John J Conlon, Dev Patel. "What Jobs Come to Mind? Stereotypes About Fields of S...	<1%
216	Publication	Jorge M. Fernandes, Javier Martínez-Cantó. "Democracy at the Crossroads – Chall...	<1%
217	Publication	José Azar, Ioana Marinescu. "Monopsony power in the labor market", Elsevier BV, ...	<1%
218	Publication	Kevin M. Williams, Teny Maghakian Shapiro. "Academic achievement across the d...	<1%
219	Publication	Komatsu, Toshiaki. "Essays on the Production Network and Labor Market", The U...	<1%
220	Publication	Luis F. Munguía Corella. "Minimum wages in monopsonistic labor markets", IZA J...	<1%

221	Publication	Lukas Hensel, Tsegay Tekleselassie, Marc Witte. "Formalized employee search an...	<1%
222	Publication	Markus A Wolf. "Persistent or temporary? Effects of social assistance benefit sanc...	<1%
223	Publication	Nicole Gürtzgen, André Diegmann (né Nolte), Laura Pohlan, Gerard J. van den Ber...	<1%
224	Publication	Paulo Somaini, Frank A. Wolak. "An Algorithm to Estimate the Two-Way Fixed Effe...	<1%
225	Publication	Philip Brownell. "Christianity and Gestalt Therapy - The Presence of God in Huma...	<1%
226	Publication	R. Jason Faberman, Guido Menzio. "Evidence on the relationship between recruiti...	<1%
227	Publication	Rania Gihleb, Osea Giuntella, Jian Qi Tan. "The impact of right-to-work laws on lo...	<1%
228	Publication	Raymundo M. Campos-Vazquez, Victor Delgado, Alexis Rodas. "The effects of a pl...	<1%
229	Publication	Springer, Lillian. "Essays on the Economics of Women's Reproductive Health.", Un...	<1%
230	Publication	Suresh Naidu, Yaw Nyarko, Shing-Yi Wang. "Monopsony Power in Migrant Labor ...	<1%
231	Publication	Thomas Le Barbanchon, Johannes Schmieder, Andrea Weber. "Job search, unempl...	<1%
232	Publication	Zen, Kai. "Essays on Information Frictions in Labor Markets", University of Califor...	<1%
233	Publication	Zeng, Yilei. "Learning Social Sequential Decision Making in Online Games", Univer...	<1%
234	Internet	amenjalal.com	<1%

235	Internet	ceistorvergata.it	<1%
236	Internet	claudioluccioletti.com	<1%
237	Internet	durham-repository.worktribe.com	<1%
238	Internet	ebin.pub	<1%
239	Internet	epub.uni-regensburg.de	<1%
240	Internet	hdl.handle.net	<1%
241	Internet	hjihao.github.io	<1%
242	Travaux envoyés	ip-paris on 2024-10-10	<1%
243	Travaux envoyés	ip-paris on 2025-05-14	<1%
244	Internet	nlsinfo.org	<1%
245	Internet	openknowledge.worldbank.org	<1%
246	Internet	publications.ut-capitole.fr	<1%
247	Internet	pure.uvt.nl	<1%
248	Internet	raw.githubusercontent.com	<1%

249	Internet	refubium.fu-berlin.de	<1%
250	Internet	repec.cepr.org	<1%
251	Internet	repository.essex.ac.uk	<1%
252	Internet	res.mdpi.com	<1%
253	Internet	researchonline.lse.ac.uk	<1%
254	Internet	senate.universityofcalifornia.edu	<1%
255	Internet	thedocs.worldbank.org	<1%
256	Internet	www.alexandria.unisg.ch	<1%
257	Internet	www.cbr.cam.ac.uk	<1%
258	Internet	www.centraalplanbureau.nl	<1%
259	Internet	www.clevelandfed.org	<1%
260	Internet	www.doria.fi	<1%
261	Internet	www.eea-esem.com	<1%
262	Internet	www.hoover.org	<1%

263	Internet	www.jaerimchoi.com	<1%
264	Internet	www.newedinburgh.ca	<1%
265	Internet	www.newyorkfed.org	<1%
266	Publication	Anna Sokolova, Todd Sorensen. "Monopsony in Labor Markets: A Meta-Analysis", ...	<1%
267	Publication	Chen, Dong. "Three essays in corporate governance", Proquest, 20111004	<1%
268	Publication	Elisa Guglielminetti, Rafael Lalive, Philippe Ruh, Etienne Wasmer. "Job search wit...	<1%
269	Publication	Jonathan Lamb, Justin Grana, Nicholas O'Donoghue. "The Benefits of Fractionati...	<1%
270	Publication	Kim, Jaeho. "Knowledge Strategies and the Direction of Innovation in Emerging T...	<1%
271	Publication	Maria De Paola, Michela Ponzio, Vincenzo Scoppa. "Gender differences in the prop...	<1%
272	Publication	Monica Langella, Alan Manning. "'Marshall Lecture 2020: The Measure of Monops...	<1%
273	Publication	Niklas Engbom, Christian Moser. "Earnings Inequality and the Minimum Wage: Ev...	<1%
274	Publication	Patrick Arni, Gerard J. van den Berg, Rafael Lalive. "Treatment Versus Regime Effe...	<1%
275	Publication	Peter Kuhn, Kailing Shen. "What Happens When Employers Can No Longer Discri...	<1%
276	Publication	Piera Bello, Alessandra Casarico, Debora Nozzay. "Research Similarity and Wome...	<1%

277	Publication	Robert Akerlof, Hongyi Li, Jonathan Yeo. "Ruling the roost: the vicious circle and t...	<1%
278	Publication	Robert Tucker Omberg. "Estimating the elasticity of labor supply to a firm: result...	<1%
279	Publication	Shouping Hu, Gregory C. Wolniak. "Initial Evidence on the Influence of College St...	<1%
280	Publication	Sosinskiy, Denis. "Essays on Minimum Wage Effects and Monopsony", University ...	<1%
281	Publication	Tommaso Giommoni. "Exposure to corruption and political participation: Evidenc...	<1%
282	Travaux envoyés	ip-paris on 2025-07-11	<1%
283	Publication	"Labor market monopsony: fundamentals and frontiers", Elsevier BV, 2025	<1%
284	Publication	Andreas Born, Eva Ranehill, Anna Sandberg. "Gender and Willingness to Lead: Do...	<1%
285	Publication	José Azar, Emiliano Huet-Vaughn, Ioana Marinescu, Bledi Taska, Till von Wachter. ...	<1%
286	Publication	Melanie Arntz, Ralf A. Wilke. "Unemployment Duration in Germany: Individual an...	<1%
287	Publication	Nguyen, Thi Ngoc Anh. "Three Essays on Political Economy.", University of Califor...	<1%
288	Publication	Ren, Kevin. "Essays in Financial Economics and Macroeconomics.", Northwestern ...	<1%
289	Publication	Ronald Bachmann, Gökay Demir, Hanna Frings. "Labor Market Polarization, Job T...	<1%
290	Publication	Ryo Kambayashi, Kohei Kawaguchi, Suguru Otani. "Decomposing Recruitment Ela...	<1%

291	Publication	
Tsuyoshi Goto, Atsushi Yamagishi. "Wage spillovers across sectors: Evidence from..."		<1%
292	Publication	
Wohl, Emma Rose. "Essays in Labor Economics.", University of California, Irvine		<1%
293	Publication	
Yim, Lokting. "EssaysOnLaborAndEducationEconomics", Purdue University, 2023		<1%



INSTITUT
POLYTECHNIQUE
DE PARIS



NNT : 20XXIPPAXXXX

Novel empirical evidence on job search

Thèse de doctorat de l'Institut Polytechnique de Paris
préparée à l'École nationale de la statistique et de l'administration économique

École doctorale n°626 École doctorale de l'Institut Polytechnique de Paris (EDIPP)
Spécialité de doctorat : Économie

Thèse présentée et soutenue à Palaiseau, le 2 novembre 2026, par

ADELAIDE FABBI

Composition du Jury :

Natalia Zinovyeva
Associate Professor, University of Warwick

Rapporteure

José de Sousa
Professeur des universités, Université Paris Panthéon-Assas
(LEMMA)

Rapporteur

Arne Uhlendorff
Directeur de recherche, CNRS (CREST)

Examineur

Maxime To
Économiste, Institut des Politiques Publiques (IPP)

Examineur

Roland Rathelot
Professeur, ENSAE – Institut Polytechnique de Paris (CREST)

Directeur de thèse

Sara Signorelli
Assistant Professor, École polytechnique – CREST (IP Paris)

Co-directrice de thèse

Thèse de doctorat

Summary

This thesis collects three independent essays in empirical labour economics. Each takes up a distinct problem in how labour markets allocate and select workers, and together they draw on three very different empirical settings and methodologies.

The first chapter investigates how an occupation-wide wage increase affects the number and quality of job applications and subsequent hiring. It exploits data from Extracadabra, an online platform for hospitality work in France, leveraging the short-term vacancies in waiting-staff positions where the default hourly wage was raised while leaving three neighbouring occupations unchanged. This setting offers a rare opportunity to observe at the same time job characteristics (wages as well as non-wage amenities), applications, applicant characteristics, and hiring decisions. A difference-in-differences design shows that the higher wage substantially increased the number of applications, pointing to a sizeable labour supply response, while mean applicant quality rose only modestly, with the improvement concentrated at the bottom of the quality distribution. Hiring increased as well. Because a higher wage in one occupation may divert applications from the others, the chapter develops a structural model of application and hiring decisions that separates the equilibrium response in the treated occupation from the spillovers to neighbouring ones. The decomposition attributes only a small share of the application effect to displacement and finds the corresponding quality spillover negligible, so that the reduced-form estimates closely approximate the equilibrium responses to the wage increase.

The second chapter turns to how public employment services are delivered. Job search assistance reaches the unemployed largely through mandatory meetings between caseworkers and job-seekers, yet what makes a given meeting effective remains something of a black box. The chapter examines one observable but overlooked feature of these meetings, the hour of the day at which they are scheduled, and asks whether such an apparently trivial choice in the operational design of a public service can affect both the effectiveness of an active labour market policy and the distribution of its benefits. Drawing on Swiss unemployment insurance registers with precisely timed meetings, it identifies the effect from within-spell variation in the meeting hour for the same job-seeker. Timing appears to matter,

though only modestly on average, and the more telling result lies in its unevenness: the early-morning advantage is considerably more pronounced for women than for men, and it is specific to that slot alone. An administrative choice made with no distributive intent thus concentrates the gains from the most effective slot among women, suggesting that the operational design of a public service shapes not only whether it works but for whom.

The third chapter studies a reform of academic hiring committees at an Italian university that paired a gender quota with the random assignment of evaluators to committees. The lottery generates quasi-exogenous variation in committee gender composition, which is exploited to identify its effect on candidates' outcomes. Combining committee records, the national registry of academic personnel, and an open bibliographic database, the chapter assembles an original dataset of hiring competitions spanning a decade, over which the reform markedly raised the average share of female evaluators. Exploiting the random assignment, it finds no evidence that committee gender balance affects a woman's probability of being selected; the estimates rule out sizeable effects but not more modest ones. A difference-in-differences design comparing fields differentially exposed to the reform points, more tentatively, in a positive direction in the most male-dominated fields.

Résumé

Cette thèse réunit trois essais indépendants d'économie du travail empirique. Chacun aborde un problème distinct de la façon dont les marchés du travail allouent et sélectionnent les travailleurs, en s'appuyant sur trois terrains et méthodologies très différents.

Le premier chapitre étudie comment une hausse de salaire à l'échelle d'un métier affecte le nombre et la qualité des candidatures ainsi que les embauches. Il exploite les données d'Extracadabra, plateforme en ligne de travail hôtelier en France, en tirant parti des missions courtes en service en salle, dont le salaire horaire par défaut a été relevé tandis que trois métiers voisins restaient inchangés. Ce cadre offre une occasion rare d'observer simultanément les caractéristiques des postes (salaires et avantages non salariaux), les candidatures, celles des candidats et les embauches. Une stratégie de doubles différences montre que la hausse de salaire a fortement accru le nombre de candidatures, traduisant une réponse marquée de l'offre de travail, tandis que la qualité moyenne des candidats n'a que légèrement augmenté, l'amélioration se concentrant au bas de la distribution. Les embauches ont aussi progressé. Une hausse de salaire dans un métier pouvant détourner des candidatures des autres, le chapitre développe un modèle structurel des décisions de candidature et d'embauche séparant la réponse d'équilibre dans le métier traité des reports vers les voisins. La décomposition n'attribue qu'une faible part de l'effet au report et juge négligeable le report de qualité, de sorte que les estimations en forme réduite approchent les réponses d'équilibre.

Le deuxième chapitre s'intéresse à la mise en œuvre des services publics de l'emploi. L'accompagnement passe largement par des entretiens obligatoires entre conseillers et demandeurs d'emploi, mais ce qui rend un entretien efficace demeure une boîte noire. Le chapitre examine une caractéristique observable mais négligée de ces entretiens, l'heure à laquelle ils sont programmés, et demande si un choix aussi anodin peut affecter à la fois l'efficacité d'une politique active et la répartition de ses bénéfices. S'appuyant sur les registres suisses de l'assurance chômage aux entretiens horodatés avec précision, il identifie l'effet à partir de la variation de l'heure de l'entretien, pour un même demandeur, au sein d'un épisode de chômage. L'heure semble compter, quoique modestement en moyenne, et le ré-

sultat le plus parlant tient à son caractère inégal : l'avantage du tôt-matin est nettement plus prononcé pour les femmes que pour les hommes, et il est propre à ce seul créneau. Un choix administratif dénué d'intention distributive concentre ainsi les gains du créneau le plus efficace sur les femmes, suggérant que la conception opérationnelle d'un service public façonne non seulement s'il fonctionne, mais pour qui.

Le troisième chapitre étudie une réforme des comités de recrutement académique dans une université italienne, qui a associé un quota de genre à l'affectation aléatoire des évaluateurs. Le tirage au sort génère une variation quasi exogène de la composition par genre, exploitée pour identifier son effet sur les résultats des candidates. En combinant procès-verbaux, registre national du personnel académique et base bibliographique ouverte, le chapitre constitue une base de données originale de concours sur une décennie, durant laquelle la réforme a nettement accru la part moyenne d'évaluatrices. En exploitant l'affectation aléatoire, il ne trouve aucune preuve que l'équilibre de genre des comités affecte la probabilité qu'une femme soit sélectionnée ; les estimations excluent des effets importants, mais pas plus modestes. Une stratégie de doubles différences comparant des disciplines inégalement exposées suggère, plus prudemment, un effet positif dans les disciplines les plus masculines.

Contents

Summary	i
Résumé	iii
General Introduction	1
1 The Wage Elasticity of Job Applications	11
1.1 Introduction	12
1.2 Institutional context and data	14
1.2.1 The online platform	14
1.2.2 Data sources	15
1.2.3 Sample and summary statistics	16
1.3 Conceptual framework	22
1.3.1 Environment	22
1.3.2 Application decisions	23
1.3.3 Equilibrium competition	24
1.3.4 Observable outcomes	27
1.4 The impact of wages on applications	29
1.4.1 Empirical strategy	29
1.4.2 Defining quality	30
1.4.3 Results	31
1.5 The impact of wages on hires	36
1.6 Structural decomposition	38
1.6.1 Parameters and moments	38
1.6.2 Identification	41
1.6.3 GMM estimator	42
1.6.4 Decomposition	42
1.6.5 Estimates and decomposition	43
1.7 Conclusion	46
Appendix 1.A Screenshots of the platform	51
1.A.1 Firm user experience	51

1.A.2	Worker user experience	52
Appendix 1.B	Estimating quality	53
Appendix 1.C	Conceptual framework: derivations and equilibrium	58
1.C.1	Worker problem	58
1.C.2	Application threshold	59
1.C.3	Worker policy functions	60
1.C.4	Exact Poisson-logit hiring operator	60
1.C.5	Mean-field approximation	61
1.C.6	Proof of Proposition 1.1	63
1.C.7	Applicant-pool outcomes	64
Appendix 1.D	Structural estimation and decomposition	64
1.D.1	Empirical moment construction	65
1.D.2	Formal identification	68
1.D.3	GMM moments	70
1.D.4	Weighting matrix and specification test	71
1.D.5	Structural decomposition	72
1.D.6	Numerical solution and inference	73
2	Caseworkers' Meetings with Job-seekers	75
2.1	Introduction	77
2.2	Literature review	79
2.3	Data	81
2.3.1	Institutional setting	81
2.4	Empirical strategy	83
2.4.1	Fixed-effects specifications	83
2.4.2	Identification and threats to validity	87
2.5	Main results: average effects of meeting timing	89
2.5.1	Average effects on the full sample	89
2.5.2	Effects across all time slots	92
2.6	Heterogeneity analysis	94
2.6.1	The gender interaction and its specificity	94
2.6.2	Other dimensions of vulnerability	95
2.6.3	Selection into the morning slot by gender	97
2.7	Robustness	100
2.7.1	The morning effect across single- and multi-spell individuals	100
2.7.2	Placebo tests on non-realised meetings	102
2.8	Discussion	106
2.8.1	Mechanism and interpretation	106
2.8.2	Policy implications	107

CONTENTS

vii

2.8.3	Limitations	108
2.9	Conclusion	109
Appendix 2.A	Institutional background	112
Appendix 2.B	Baseline OLS analysis	113
Appendix 2.C	Considered and discarded designs	121
Appendix 2.D	Intermediate outcomes and the duration contamination	121
Appendix 2.E	Appendix tables	122
3	Gender-Balanced Committees	135
3.1	Introduction	135
3.1.1	Composition of the committee	138
3.1.2	Gender quotas	140
3.2	Institutional Background and the 2019 Reform	143
3.2.1	The Gender Equality Reform of July 2019	144
3.3	Data	146
3.3.1	Limitations of the current dataset	147
3.3.2	Summary Statistics	149
3.4	Does Committee Composition Affect Women's Success?	152
3.4.1	Empirical Strategy	152
3.4.2	Main results	154
3.4.3	Statistical power and minimum detectable effects	156
3.5	Did the Reform Improve Outcomes for Women?	157
3.5.1	A continuous measure of treatment intensity	157
3.5.2	The reform's differential effect across fields	161
3.6	Conclusion	163
Appendix 3.A	Regulatory History and National Rules	170
3.A.1	Regulatory Changes and Ongoing Reform	170
3.A.2	A National Survey on Committee Composition Rules	171
3.A.3	Summary Statistics and descriptive	172
3.A.4	Post Reform Analysis	175
3.A.5	Pre-post Reform Analysis	189

CONTENTS

General Introduction

This thesis brings new empirical evidence to three moments of the job search process: how workers respond to what a job pays, how public assistance shapes their search, and how candidates are selected.

The three chapters of this thesis share several features. (i) They study the frictions and interventions that shape how workers and jobs come together, looking at both sides of the match: the first two chapters focus on the job seeker, while the third turns to the selection side. (ii) They draw on rich individual-level data, whether administrative or platform-based. (iii) They rely on credible research designs to isolate causal effects, exploiting quasi-experimental variation, random assignment, and difference-in-differences with continuous treatment intensity. (iv) Each speaks to a concrete policy lever, from the design of platform and public employment services to diversity policies in hiring, and each pays attention to the distributional dimension of these questions.

The first chapter asks how workers respond when a job's pay rises: not only whether more of them apply, but whether the applicants are any better. This connects to a long literature **estimating the elasticity of labour supply to jobs and firms** (Banfi and Villena-Roldán, 2019; Marinescu and Wolthoff, 2020; Manning, 2021; Card, 2022).

Wages are not the only aspect under the lens: a large part of the field asks how much weight non-wage amenities carry in the same choice (Mas and Pallais, 2017; Maestas et al., 2023). We observe several of these, but here they enter as controls rather than as the object of study. Another strand looks at how long vacancies stay open as wages change (Bassier et al., 2025). Answering these questions, however, is empirically challenging. Researchers rarely observe the entire market in which workers and firms interact. Estimating labour supply responses requires following the full recruitment process, from the jobs available and the workers who are exposed to them, to the applications they submit and the hiring decisions that follow. Such data are exceptionally rare.

The Extracadabra data provided a unique opportunity to study these questions. This chapter, joint with Louise Demoury and Roland Rathelot, grew out of a collaboration with the platform, an online marketplace for short-term hospitality

work in France, mostly in the Île-de-France region (and above all in Paris). We built the dataset ourselves from its records, following all three stages, notification, application, and hire, on the same jobs. For every job we see who the platform notified, including workers who never applied, each application and its timestamp, whether and when the employer ignored, rejected, or selected each applicant, and, for completed jobs, the hours worked and the wage bill.

The data reveal information on both sides. On the restaurant side, each job records its occupation, the venue's postcode and sector, and, for bars and restaurants, a detailed establishment category, together with the full shift schedule down to the unpaid breaks, the number of workers required, and the date and time the position was created. On the worker side, each registered worker has a profile with age, gender, postcode of residence, declared skills, and previous work experience; because the platform updates the profile as assignments are completed, we also see the worker's past jobs and the ratings employers gave them, as well as every time they connect to the app. In all, the sample spans roughly thirteen months and covers 26,297 job adverts from 1,475 firms across 2,200 venues, and 5,632 workers who together submitted 320,833 applications.

Before any analysis or regression, understanding the platform was itself part of the project: for example how its matching algorithm decides which workers are shown which jobs, but more importantly how pay is set. We started by looking at two kinds of arrangement together, permanent contracts and short "extra" shifts, and the design came into focus only as we understood the second. Extras make up most of the platform and last from a few days to a few weeks. Crucially, the platform sets a default hourly wage for each occupation that most employers simply accept, and because payment runs through the platform, the wage we observe is the wage actually paid, with no room for off-platform deals. That is what makes the setting clean. On 5 October 2023 the platform raised the waiting-staff default from 16 to 17 euros while leaving bar staff, runners, and receptionists unchanged, a dated, occupation-specific change we can treat as close to a natural experiment and study with an event-study design.

With the setting in place, we turned to measuring quality. In our case, it did not mean an absolute judgment of a worker: we mean how good a fit an applicant looks for a given job. We build it from employers' revealed preferences, a worker-job-specific score for whether an applicant, conditional on applying, is hired, so the same worker can score differently for different jobs. The score draws on things like whether the worker has been approved by that employer before, distance from home to the workplace, average rating, experience in the occupation and type of establishment, and how recently they were active on the app.

A last concern is hiring itself. We measure it as the number of workers a job

hires from among its applicants, but unlike an application, which is the worker's decision alone, a hire is a joint decision of worker and employer. The effect we estimate on hiring therefore combines the extra applications a higher wage brings with any change in how selectively employers choose. **Because the default wage is not binding, these are intention-to-treat effects.**

Applications rose by approximately 23 per cent relative to the comparison occupations, an elasticity of about 3.5. Mean applicant quality rose by approximately 2 per cent of its pre-treatment mean, an elasticity of about 0.3. The minimum quality of the applicant pool increased more strongly, with an elasticity of approximately 1.7, **while the effect on maximum quality is small and statistically insignificant, the confidence interval ruling out an elasticity above approximately 0.4.** The number of hires increased by about 11 per cent, an elasticity of 1.8. The wage increase therefore attracted substantially more applicants and modestly improved the applicant pool, primarily by raising its lower tail. The reduced form leaves a further problem. Raising the wage in one occupation does not leave the others alone: workers can shift towards the better-paid role, so the treatment reaches the control occupations too, a violation of SUTVA. To deal with it, we built a model of how workers apply and employers hire, calibrated to the platform, to act as support. It splits the reduced-form comparison into the equilibrium effect on the occupation whose wage rose and the applications merely displaced from the others. The displacement turns out to be small: spillovers account for only about 6 per cent of the measured effect, so the simple estimates are close to the true equilibrium ones. This chapter also bears on an older question: how much power employers have to set wages. The idea that firms may face an upward-sloping labour supply curve goes back to Robinson, 1933, who first described employers with some degree of monopsony power. Although this perspective was long overshadowed by the competitive model, it has returned to the centre of labour economics, with a growing literature asking how far real labour markets depart from perfect competition (Manning, 2021; Card, 2022; Sokolova and Sorensen, 2021). Our estimate adds one number to this debate: how strongly workers redirect their search when one job's pay rises.

The second chapter stays with job search but turns to the state's attempt to help it along. Public employment services deliver most of that support through one instrument, the mandatory meeting between a caseworker and an unemployed job-seeker. A substantial literature establishes that these meetings are a productive input and identifies who should conduct them and how often (Behncke et al., 2010; Schiprowski, 2020; **Van den Berg et al., 2012**; Berg and Klaauw, 2006; Maibom et al., 2017; Böheim et al., 2022). The meeting itself, however, remains a black box: administrative registers record that an interaction took place, not what happened

inside it.

The natural way in is meeting duration, the time a caseworker devotes to a job-seeker, and this is where the project started. The register does record it, but after cleaning the durations concentrate at a small number of mass points, with spikes at 30, 45, and 60 minutes. They are reliably recorded, but only at rounded values, which leaves too little variation within each group to identify a marginal effect of meeting length. Duration therefore enters the analysis as a control rather than as a treatment.

One dimension of the interaction the register does capture, and that the literature has not used, is the hour at which the meeting is scheduled. There is reason not to presume it neutral: students score higher in morning exams (Pope, 2016), and recruiters on a Swiss public employment platform screen applications less carefully before lunch and at the end of the day, amplifying discrimination against minority candidates (Hangartner et al., 2021). The question is therefore broader than whether early meetings beat late ones. It is whether apparently trivial choices in the operational design of a public service, such as the distribution of appointments over the working day, affect both the effectiveness of an active labour market policy and the distribution of its benefits.

The data are the Swiss unemployment insurance registers covering all registered unemployed between 2001 and 2018, built from two linked systems, one for placement and labour market statistics and one for benefit payments. The full register holds 20,425,634 meetings, each with an exact timestamp and a recorded duration, alongside 247 variables on job-seekers, their employment histories, caseworkers, and the meetings themselves. Extensive cleaning was required: we keep only meetings that took place, at plausible hours and with plausible recorded durations, which removes the pro-forma entries massed at the administrative default, leaving an analysis sample of 11,745,735 meetings. Meetings that were postponed, cancelled, or not attended are set aside rather than discarded, and serve as the basis for the placebo tests.

The setting is not a natural experiment, so identification rests on within-individual comparisons. We compare meetings held at different hours by the same job-seeker, within the same unemployment spell and at the same stage of it, with individual-by-spell and caseworker fixed effects. This is demanding: a spell contributes only if it contains more than one meeting at distinct hours, which brings the estimation sample to 8,832,592 meetings. Meeting times are set by caseworkers and accepted by job-seekers rather than self-selected, so the residual variation plausibly reflects caseworker-side scheduling constraints, agenda availability, substitutions, office reorganisation, rather than job-seeker choice. We probe this interpretation with a battery of placebo tests on meetings that were scheduled but never took place.

Meeting timing matters, but modestly. An early-morning meeting, held between 7 and 8 a.m., is associated with a 0.22 percentage point higher probability of employment six months later and 0.019 additional months of employment over the following half-year, relative to the 9 to 10 a.m. reference slot. This modest average conceals substantial heterogeneity. The association is roughly twice as strong for women as for men, 0.32 against 0.15 percentage points on six-month employment, and it is specific to the early-morning slot: interactions between gender and every other time slot are precise zeros. Women do not simply do better at all meetings. This specificity makes a data-mining explanation less likely, since a chance rejection would not single out one slot, and the pattern is stable across alternative clustering, the exclusion of the rare evening slot, and an alternative outcome.

We read the gender result as an equity finding rather than as evidence of discrimination, and the estimates as reduced-form, since the data do not identify the mechanism behind the timing effect. The contribution of the chapter is to open one dimension of this black box: intra-day scheduling, a margin of policy delivery that is observable, administratively set, and cheap to change, yet largely overlooked. The average return is modest; what stands out is its distribution. An administrative choice with no distributive intent concentrates the benefits of the most effective slot among women, so the operational design of the service shapes not only whether the policy works but for whom. We are currently extending the analysis to the updated register data, with short-horizon outcomes starting one month after the meeting; the intensity margin, how much time a caseworker devotes to each job-seeker, remains the natural complement to timing and is out of reach until registers record durations more finely.

Applications and job-search assistance are only one side of the market; the other is who gets chosen. The third chapter turns to the people who do the choosing, and asks whether the gender composition of academic hiring committees changes a woman's chances of being selected.

Evidence on this already exists, but it does not travel well. Studies in different places have reached opposite conclusions: in Spain more women on a committee left women's chances broadly unchanged, while in Italy an extra female evaluator, if anything, widened the gap (Bagues et al., 2017); a French quota raised the share of women on committees but lowered the chance a woman was ranked first (Deschamps, 2024); and applications evaluated more favourably when they come from women in Germany look gender-neutral in Italy (Gërkhani et al., 2025). Recent work suggests this is not just noise: the effect of committee composition depends on the surrounding norms and can even change sign over time (Baron et al., 2024).

There is a further reason to look at Italy now. Gender-balance requirements

for scientific and hiring committees have spread across Europe, often set at a minimum of one third of each gender, with many universities adopting their own internal guidelines (E.C., 2018). In Italy these requirements are largely left to individual universities and vary widely: of the forty-one universities that answered our 2023 survey, only seven pair a one-third quota with a random draw, so Pisa's arrangement is uncommon.

Pisa introduced it in 2019, a binding one-third quota enforced by a lottery, where before there had been no gender-balance requirement at all. Pisa's reform therefore provides an opportunity to evaluate not only whether committee composition matters, but also whether the policy designed to change it achieved its intended objective.

The dataset required combining three sources that had never been brought together for this purpose: the University of Pisa's published competition records, which give the exact composition of every committee and the list of candidates; the national CINECA registry, for each evaluator's age, rank, and field, and for the pre-reform gender balance of each discipline; and the OpenAlex database, scraped for the research output of every evaluator and candidate. Names had to be matched to the right scholarly profile within each disciplinary sector, and publication records reconstructed as they stood at the date of each competition. The result covers 789 competitions and 3,150 candidates between 2015 and 2024.

Pisa's reform gives us two sources of variation, and the chapter uses both. Because committee members are drawn by lot, the gender composition of a committee is quasi-random, which lets us estimate what one more woman on a committee does to a candidate's chance of winning, at the level of an individual hiring, where random assignment of evaluators is otherwise rare. The reform itself gives the second: fields that started further from gender parity had further to move when the one-third quota came in, so comparing candidates' outcomes before and after across fields more or less exposed to the reform asks a different question, not what one extra woman on a committee does, but what the whole policy did. Bringing these two together in a single institution is itself much of what the chapter contributes.

The two designs give different answers, and neither is conclusive on its own. Using the lottery, we find no evidence that the gender composition of a committee affects a woman's probability of being selected. The estimate is cleanly identified but imprecise: it rules out large marginal effects and cannot exclude smaller ones, so this is an absence of evidence of a large effect rather than evidence of no effect. The difference-in-differences goes the other way, suggestive in sign but weakly identified, resting on a short and uneven pre-reform window that does not allow a convincing test of parallel trends. At the average field the reform

REFERENCES

7

leaves women's outcomes unchanged; the effect appears only as fields become more male-dominated, where it points, tentatively, in a positive direction. What the two share is their direction: neither produces the sizeable pro-female effect a mechanical reading of gender balance would predict, and the only movement we detect is confined to the fields that started furthest from parity.

Both limitations point the same way. The lottery estimate is imprecise because a single university yields only so many competitions: halving the smallest effect we could detect would take roughly four times as many, which means more years of data or other universities that also draw their committees by lot. The pre-reform window is short and still being completed, which is why we treat the difference-in-differences as exploratory. And the evidence comes from one institution, so whether it carries over to universities with different disciplinary mixes or hiring cultures is not something our data can settle.

These limitations naturally point to future work: more competitions, and ideally more universities with a comparable mechanism, would sharpen what we can say. More broadly, the chapter contributes to the growing literature on gender quotas in academic recruitment by showing that raising women's representation on evaluation committees does not by itself produce immediate changes in hiring outcomes.

Taken together, the three chapters study different stages of labour market matching. The first asks how workers respond to economic incentives, the second how public institutions shape search behaviour, and the third how employers evaluate candidates.

References

Bagues, Manuel, Mauro Sylos-Labini, and Natalia Zinovyeva (Apr. 2017). "Does the Gender Composition of Scientific Committees Matter?" In: *American Economic Review* 107.4, pp. 1207–38. DOI: 10.1257/aer.20151211. URL: <https://www.aeaweb.org/articles?id=10.1257/aer.20151211>.

Banfi, Stefano and Benjamín Villena-Roldán (July 2019). "Do High-Wage Jobs Attract More Applicants? Directed Search Evidence from the Online Labor Market". In: *Journal of Labor Economics* 37.3, pp. 715–746. ISSN: 0734-306X. DOI: 10.1086/702627. URL: <https://www.journals.uchicago.edu/doi/10.1086/702627> (visited on 03/29/2024).

Baron, Justus, Bernhard Ganglmair, Nicola Persico, Timothy Simcoe, and Emanuele Tarantino (2024). "Representation is not sufficient for selecting gender diversity". In: *Research Policy* 53.6, p. 104994.

- 25 Bassier, Ihsaan, Alan Manning, and Barbara Petrongolo (2025). "Vacancy Duration and Wages". In: *Review of Economics and Statistics*, pp. 1–28. DOI: [10.1162/rest_a_01580](https://doi.org/10.1162/rest_a_01580).
- 12 Behncke, Stefanie, Markus Frölich, and Michael Lechner (2010). "Unemployed and their caseworkers: should they be friends or foes?" In: *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 173.1, pp. 67–92. DOI: <https://doi.org/10.1111/j.1467-985X.2009.00600.x>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-985X.2009.00600.x>.
- 82 7 Berg, Gerard van den and Bas van der Klaauw (2006). "Counseling and Monitoring of Unemployed Workers: Theory and Evidence from a Controlled Social Experiment." In: *International Economic Review* 47.3, pp. 895–936. URL: <https://EconPapers.repec.org/RePEc:ier:iecrev:v:47:y:2006:i:3:p:895-936>.
- 15 60 Böhme, René, Rainer Eppel, Helmut Mahringer, et al. (2022). *More Caseworkers Shorten Unemployment Durations and Save Costs. Results from a Field Experiment in an Austrian Public Employment Office*. Tech. rep. WIFO.
- 42 8 Card, David (Apr. 2022). "Who Set Your Wage?" en. In: *American Economic Review* 112.4, pp. 1075–1090. ISSN: 0002-8282. DOI: [10.1257/aer.112.4.1075](https://doi.org/10.1257/aer.112.4.1075). URL: <https://www.aeaweb.org/articles?id=10.1257/aer.112.4.1075> (visited on 08/03/2023).
- 2 Deschamps, Pierre (2024). "Gender quotas in hiring committees: A boon or a bane for women?" In: *Management Science* 70.11, pp. 7486–7505.
- 117 E.C. (2018). *Guidance to facilitate the implementation of targets to promote gender equality in research and innovation*. Tech. rep. European Commission - DG for Research and Innovation.
- 21 Gërzhani, Klarita, Nevena Kulic, Alessandra Rusconi, and Heike Solga (2025). "Gender bias in evaluating assistant professorship applicants? Evidence from harmonized survey experiments in Germany and Italy". In: *Social Science Research* 126, p. 103113.
- 10 Hangartner, Dominik, Daniel Kopp, and Michael Siegenthaler (2021). "Monitoring hiring discrimination through online recruitment platforms". In: *Nature* 589.7843, pp. 572–576.
- 26 Maestas, Nicole, Kathleen J Mullen, David Powell, Till Von Wachter, and Jeffrey B Wenger (2023). "The value of working conditions in the United States and implications for the structure of wages". In: *American Economic Review* 113.7, pp. 2007–2047.
- 37 7 Maibom, Jonas, Michael Rosholm, and Michael Svarer (2017). "Experimental Evidence on the Effects of Early Meetings and Activation". In: *The Scandinavian Journal of Economics* 119.3, pp. 541–570. DOI: <https://doi.org/10.1111/sjoe.12180>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/sjoe.12180>.

REFERENCES

9

12180. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/sjoe.12180>.

Manning, Alan (Jan. 2021). "Monopsony in Labor Markets: A Review". en. In: *ILR Review* 74.1, pp. 3–26. ISSN: 0019-7939. DOI: 10.1177/0019793920922499. URL: <https://doi.org/10.1177/0019793920922499> (visited on 08/21/2023).

Marinescu, Ioana and Ronald Wolthoff (Apr. 2020). "Opening the Black Box of the Matching Function: The Power of Words". In: *Journal of Labor Economics* 38.2, pp. 535–568. ISSN: 0734-306X. DOI: 10.1086/705903. URL: <https://www.journals.uchicago.edu/doi/abs/10.1086/705903> (visited on 08/03/2023).

Mas, Alexandre and Amanda Pallais (2017). "Valuing alternative work arrangements". In: *American Economic Review* 107.12, pp. 3722–3759.

Pope, Nolan G (2016). "How the time of day affects productivity: Evidence from school schedules". In: *Review of Economics and Statistics* 98.1, pp. 1–11.

Robinson, Joan (1933). *The economics of imperfect competition*. Springer.

Schiprowski, Amelie (2020). "The Role of Caseworkers in Unemployment Insurance: Evidence from Unplanned Absences". In: *Journal of Labor Economics* 38.4, pp. 1189–1225. DOI: 10.1086/706092. eprint: <https://doi.org/10.1086/706092>. URL: <https://doi.org/10.1086/706092>.

Sokolova, Anna and Todd Sorensen (Jan. 2021). "Monopsony in Labor Markets: A Meta-Analysis". en. In: *ILR Review* 74.1, pp. 27–55. ISSN: 0019-7939. DOI: 10.1177/0019793920965562. URL: <https://doi.org/10.1177/0019793920965562> (visited on 08/21/2023).

Van den Berg, Gerard J., Lene Kjaersgaard, and Michael Rosholm (2012). "To Meet or Not to Meet (Your Case Worker) – that is the Question." In: *IZA Discussion Paper* 6476.

Chapter 1

The Wage Elasticity of Job Applications: Quantity and Quality Margins

Abstract

We study how an occupation-wide wage increase affects the number and quality of job applications and subsequent hiring. We use data from *Extracadabra*, an online platform for short-term hospitality work in France, which records job notifications, applications, and hiring decisions. The platform raised the default hourly wage for waiting-staff jobs from €16 to €17 in October 2023 while leaving the defaults for three neighbouring occupations unchanged. A difference-in-differences design shows that the wage increase raised applications by approximately 23%, implying a wage elasticity of about 3.5. Mean applicant quality rose by 2%, an elasticity of approximately 0.3, with the largest improvement at the bottom of the quality distribution. Hires rose by about 11%, implying an elasticity of 1.8. Because the wage change may divert applications from the comparison occupations, the usual no-interference assumption need not hold. We therefore estimate a model of application and hiring decisions that separates the equilibrium response in the treated occupation from spillovers to neighbouring occupations. The decomposition assigns about 6% of the application difference-in-differences estimate to displacement from the comparison occupations and essentially none of the quality effect to spillovers. The reduced-form estimates therefore closely approximate the equilibrium responses in the treated occupation.

This chapter is co-authored with Louise Demoury and Roland Rathelot.

1.1 Introduction

Employers in hospitality, construction, and health care often report persistent recruitment difficulties, even when unemployment remains substantial. In a frictionless labour market, wages would adjust to direct workers towards occupations with high labour demand. How strongly workers respond to such wage differences is therefore central to understanding both recruitment and the allocation of labour across occupations. Existing evidence largely concerns the number of applications or hires. Much less is known about the quality of the workers attracted by higher wages, partly because applications, applicant characteristics, and hiring decisions are rarely observed in the same data.

This article studies how an occupation-wide wage increase affects the quantity and quality of applications and subsequent hiring. We use data from *Extracadabra*, an online platform for short-term hospitality work in and around Paris. The platform sets an occupation-specific default hourly wage, which most employers adopt. On 5 October 2023, it raised the default wage for waiting-staff jobs from €16 to €17, an increase of 6.25%, while leaving the defaults for bar staff, runners, and receptionists unchanged. Comparing waiting-staff vacancies with vacancies in these three occupations before and after the change identifies the response to a change in an occupation's relative wage rather than to a wage increase at a single firm.

The platform data record the workers notified about each vacancy, their applications, and employers' hiring decisions. We use these hiring decisions and rich worker-job characteristics to construct an applicant-level quality index. The data therefore allow us to distinguish four margins: the number of applications, the mean quality of the applicant pool, its lower and upper tails, and realised hiring.

Our preferred difference-in-differences specification shows that the wage increase raised the number of applications by approximately 23% relative to the comparison occupations. Dividing the change in log applications by the change in the log wage gives an elasticity of about 3.5. Mean applicant quality rose by approximately 2% of its pre-treatment mean, implying an elasticity of about 0.3. The minimum quality in the applicant pool increased more strongly, with an elasticity of approximately 1.7. The estimated effect on maximum quality is small and statistically insignificant; the confidence interval rules out an elasticity above approximately 0.4. Finally, the number of hires increased by about 11%, implying an elasticity of 1.8. Thus, the wage increase attracted substantially more applicants and modestly improved the applicant pool, primarily by raising its lower tail.

The reduced-form contrast does not, by itself, identify the equilibrium effect on

1.1. INTRODUCTION

13

waiting-staff jobs. Workers can apply across occupations, so the higher waiting-staff wage may divert applications from the comparison occupations. This interference violates the stable unit treatment value assumption underlying a conventional difference-in-differences interpretation. We therefore develop a model of workers' application decisions and employers' hiring decisions. Workers differ in their productivity in two occupations and account for both the wage and their probability of being hired. A higher wage can attract productive workers, intensify competition, and discourage applicants with weaker hiring prospects. The model separates the equilibrium response in the treated occupation from the spillover to the comparison occupations.

Estimating the model shows that spillovers account for about 6% of the model-based application contrast: the higher waiting-staff wage reduces applications to the comparison occupations, modestly amplifying the difference-in-differences estimate. The remaining 94% reflects the equilibrium response in the treated occupation. The corresponding quality spillover is negligible. The reduced-form estimates therefore provide close approximations to the equilibrium effects of the wage increase on waiting-staff vacancies.

This article contributes to three literatures. First, it contributes to work estimating the elasticity of labour supply to firms and jobs using applications, vacancy duration, recruitment, quits, and employment flows (Ransom and Sims, 2010; Staiger et al., 2010; Matsudaira, 2014; Banfi and Villena-Roldán, 2019; Dube et al., 2020; Marinescu and Wolthoff, 2020; Bassier et al., 2022; Sokolova and Sorensen, 2021; Manning, 2021; Card, 2022; Hirsch et al., 2022; Mueller et al., 2023; Datta, 2023; Bassier et al., 2025). Banfi and Villena-Roldán, 2019 and Marinescu and Wolthoff, 2020 provide pioneering evidence about the correlation between posted wages and the number of applications received by job ads. Belot et al., 2022 and Beerli et al., 2026 study how jobseekers respond to experimentally varied wage information. Hirsch et al., 2022, Datta, 2023, Mueller et al., 2023, and Bassier et al., 2025 focus on recruitment outcomes (vacancy duration and hires). Our source of variation differs from most of the literature because the wage changes for an entire occupation. Our design allows us to capture both the response to the higher wage and the reallocation of applications across contemporaneously available jobs. To our knowledge, our paper is the first attempt to decompose this market-level effect between an effect on the affected occupation and spillovers. Another contribution is that we can measure the impact on the number of applications, application quality, and hires within the same setting.

Second, the article contributes to the smaller literature on how wages affect applicant composition. Dal Bó et al., 2013 provide experimental evidence that higher

public-sector wages attract a larger and more qualified applicant pool. Related work studies how wage information and wage floors affect applicant selection, skill requirements, and the composition of job offers (Clemens et al., 2021; Balgova et al., 2025). We observe applications, applicant characteristics, and hiring outcomes on the same platform, allowing us to distinguish changes in the number of applicants from changes throughout the quality distribution and from realised hiring.

Third, the article relates to work on wage floors, reallocation, and occupational mobility. Wage-setting policies can reallocate workers and firms across jobs, establishments, and local labour markets (Cengiz et al., 2019; Harasztosi and Lindner, 2019; Dustmann et al., 2022; Derenoncourt and Montialoux, 2021; Caldwell and Danieli, 2024; Azar et al., 2024; Derenoncourt and Weil, 2025; Giupponi et al., 2024), and can generate spillovers to untreated workers, firms, or markets (Gopalan et al., 2021; Bassier, 2022; Demir, 2024; Giupponi and Machin, 2026). Other work studies how information and search assistance affect mobility across occupations and employers (Jäger et al., 2024; Altmann et al., 2022; Belot et al., 2019; Belot et al., 2025; Belot et al., 2026; Klæui, 2026). Our contribution is to estimate an occupation-level wage response and to provide a framework for separating the response in the treated occupation from the displacement of applications elsewhere on the platform.

Section 2 describes the institutional setting and data. Section 3 develops the conceptual framework. Sections 4 and 5 present the reduced-form evidence on applications and hires. Section 6 estimates the structural model and decomposes the reduced-form effects. Section 7 concludes.

1.2 Institutional context and data

1.2.1 The online platform

Extracadabra is an online platform in the French labour market that specialises in hospitality jobs in and around Paris. Its clients are mainly restaurants, bars, hotels, and event organisers. The platform advertises three types of arrangement: standard employment contracts of fixed or indefinite duration, temporary-employment contracts, and self-employed contractor arrangements in which workers supply services through their own businesses. We focus on the last category, which is the most common on the platform and on which the richest data are available. These assignments typically last from a few hours to a few days and must be filled at short notice. In the French hospitality industry, such occasional shifts are known

1.2. INSTITUTIONAL CONTEXT AND DATA

15

as *extras*.

Extracadabra differs from a conventional job board because workers do not search the full set of posted vacancies. An employer first provides the characteristics of the position to the platform (Appendix 1.A.1). A matching algorithm then identifies a set of suitable workers and sends each of them a notification through the mobile application. Notified workers can inspect the job details and either ignore the notification, decline the opportunity, or apply (Appendix 1.A.2). The employer is notified of each application, reviews the applicant profiles, and selects one or more workers or leaves the position unfilled. Once a worker is selected, the platform charges the employer the wage bill and its commission and holds the payment in escrow.

After the job is completed, the employer can rate the worker on a scale from 1 to 5. Ratings are usually either missing or equal to 4 or 5. A worker rated 4 or above becomes an approved worker for that employer and receives priority in the matching algorithm for the employer's subsequent vacancies in any occupation.

The platform also plays an unusually direct role in wage setting. It specifies a default hourly wage for each occupation. Employers can depart from this default by asking customer service to create a client-specific wage grid. Since 19 December 2023, they have also been able to boost the wage when creating a position. Between April 2023 and May 2024, 94.6% of posted wages equalled the occupation-wide default, 4.5% equalled a client-specific default, and 0.9% used the wage-boost option. Our empirical design exploits a change in one of these occupation-specific defaults.

1.2.2 Data sources

The platform provides job-level, worker-level, and worker-job-level data.

Job-level data

For each job, we observe the occupation, workplace postcode, complete schedule (including unpaid breaks), number of workers required, and the date and time at which the position was created. We also observe the employer's identity, location, sector, and, for bars and restaurants, a detailed establishment category such as "Italian restaurant".

Worker-level data

Workers must register before applying. Their profiles contain age, gender, postcode of residence, declared skills, and previous work experience. The platform updates these profiles as workers complete assignments, so at each date we observe a worker's prior assignments and the associated employer ratings. We also observe every time a worker connects to the mobile application.

Pair-level data

A distinctive feature of the data is that they cover each stage of the matching process. For every job, we observe the workers notified by the platform, including those who do not apply. We then observe each application and its timestamp, whether and when the employer ignored, rejected, or selected the applicant, and the workers who ultimately performed the job. For completed jobs, we observe hours worked and the wage bill transferred through the platform.

1.2.3 Sample and summary statistics

The analysis sample contains jobs posted between 5 April 2023 and 4 May 2024 for waiting staff and three neighbouring occupations: bar staff, runners, and receptionists. We exclude the small minority of jobs that platform staff intermediated manually.

The sample contains 26,297 job adverts posted by 1,475 firms at 2,200 venues, where a venue is an establishment–location pair. At least one worker was hired for 47.1% of these jobs. The worker sample comprises the 5,632 workers who submitted at least one application to a sample job; together, they made 320,833 applications.

Table 1.1 reports job characteristics by occupation. Jobs last an average of 6 hours, and only 20% span more than one day. They include an average of 30 minutes of unpaid breaks and are posted 3.5 days before they begin. Employers request an average of 1.3 workers per job. Almost all jobs are in the Paris region, and 66–88% are within the city of Paris, depending on the occupation. The platform notifies an average of 774–1,335 workers per job, of whom 8.2–19.4 apply. Employers hire an average of 0.55 workers per job, and approximately half of all jobs result in at least one hire.

1.2. INSTITUTIONAL CONTEXT AND DATA

17

Table 1.1: Summary statistics for jobs

	Waiting staff	Bar staff	Receptionist	Runner	Non-waiting staff
Nb. of matches	1,335.09 (496.71)	898.50 (297.72)	1,159.87 (471.15)	773.99 (197.59)	885.07 (314.65)
Nb. of applications	13.77 (12.55)	10.61 (9.63)	19.38 (13.67)	8.15 (8.31)	10.74 (10.27)
Nb. of hires	0.55 (0.71)	0.61 (0.78)	0.59 (0.73)	0.49 (0.66)	0.57 (0.74)
At least one hire	0.46 (0.50)	0.50 (0.50)	0.48 (0.50)	0.43 (0.50)	0.48 (0.50)
Nb. of extras demanded	1.31 (1.27)	1.27 (1.30)	1.35 (1.17)	1.23 (1.30)	1.27 (1.28)
Nb. of work hours	5.75 (1.94)	5.86 (1.82)	6.08 (2.03)	6.91 (1.85)	6.24 (1.92)
Nb. of break hours	0.30 (0.61)	0.33 (0.55)	0.39 (0.49)	0.40 (0.55)	0.36 (0.55)
Time until shift start (hours)	88.91 (141.89)	58.26 (97.83)	112.43 (151.39)	111.84 (177.76)	82.28 (138.49)
Located in Paris	0.66 (0.47)	0.69 (0.46)	0.88 (0.32)	0.71 (0.45)	0.72 (0.45)
Located in IDF	0.96 (0.21)	0.91 (0.28)	0.99 (0.12)	0.85 (0.35)	0.90 (0.30)
Morning	0.14 (0.35)	0.10 (0.30)	0.23 (0.42)	0.07 (0.26)	0.11 (0.31)
Lunch	0.44 (0.50)	0.43 (0.49)	0.48 (0.50)	0.26 (0.44)	0.37 (0.48)
Afternoon	0.27 (0.44)	0.31 (0.46)	0.39 (0.49)	0.27 (0.44)	0.31 (0.46)
Evening	0.34 (0.47)	0.30 (0.46)	0.32 (0.47)	0.24 (0.43)	0.28 (0.45)
Night	0.25 (0.43)	0.33 (0.47)	0.24 (0.43)	0.57 (0.49)	0.40 (0.49)
Nb. of working days					
1	0.81	0.79	0.74	0.76	0.77
2	0.09	0.10	0.11	0.12	0.11
3	0.03	0.04	0.06	0.04	0.04
4	0.02	0.02	0.02	0.02	0.02
5	0.03	0.03	0.05	0.04	0.04
>5	0.02	0.02	0.03	0.02	0.02
Type of shift					
Week	0.64	0.55	0.60	0.56	0.56
Weekend	0.06	0.09	0.12	0.08	0.09
Week & Weekend	0.30	0.36	0.29	0.36	0.35
Share of sample	0.48	0.29	0.06	0.17	0.52
N	12,677	7,530	1,491	4,599	13,620

Notes: The table reports summary statistics for jobs posted between 5 April 2023 and 4 May 2024 in the waiting-staff, bar-staff, receptionist, and runner occupations. Standard deviations are reported in parentheses below means. No additional statistic is reported below proportions.

Figure 1.1 shows the number of jobs created each month. Postings are strongly seasonal: demand is highest in June and September and is approximately twice as high as in July and August. Workers do not observe the employer's name when applying, which helps prevent employers and workers from bypassing the platform. Instead, they see a detailed category describing the type of establishment. Figure 1.2 displays the distribution of establishment categories. The largest categories are café-restaurants, Italian restaurants, fast-food establishments, three-star hotels, and catering services.

Table 1.2 reports applicant characteristics in Panel A and employer characteristics in Panel B. Because workers can apply in several occupations, the occupation-specific columns are not mutually exclusive. Waiting-staff applicants are 28 years old on average, and 54% are women. Almost all (99%) live in Île-de-France, and 74% live within the city of Paris. Their homes are an average of 17 kilometres from the job: 29% live in the same *département* as the workplace and 1% in the same postcode. Applicants are highly active on the platform. Before receiving a job notification, 56% had connected within the previous hour and 86% within the previous 12 hours. Their average rating is 4.5, and 0.8% have no previous rating.


Applicant characteristics are broadly similar across occupations. Gender is a notable exception: receptionist applicants are more likely to be women, bar-staff and runner applicants are more likely to be men, and waiting-staff applicants lie between these groups.

Panel B describes the 1,475 firms in the sample. Because firms can post jobs in several occupations, they may also appear in several columns. Firms posting waiting-staff jobs operate an average of 1.5 venues, compared with approximately 1.2 among firms posting in the other three occupations. Conditional on posting in an occupation, firms create an average of 12 adverts during the sample period.

1.2. INSTITUTIONAL CONTEXT AND DATA


19

Figure 1.1: Monthly job postings



Notes: The figure reports monthly job postings for waiting staff, bar staff, receptionists, and runners from 5 April 2023 to 4 May 2024.

Figure 1.2: Job postings by establishment category



Notes: The figure reports the distribution of job postings across establishment categories for waiting staff, bar staff, receptionists, and runners from 5 April 2023 to 4 May 2024. Categories with fewer than 50 postings are grouped as "Other".

CHAPTER 1. THE WAGE ELASTICITY OF JOB APPLICATIONS

Table 1.2: Summary statistics of workers and firms

	Waiting staff	Bar staff	Receptionist	Runner	Non-waiting staff
Panel A: applicant characteristics					
Age	27.96 (6.52)	28.55 (6.42)	26.98 (5.41)	29.06 (6.89)	28.37 (6.40)
Female	0.54 (0.50)	0.36 (0.48)	0.71 (0.45)	0.39 (0.49)	0.44 (0.50)
Avg. rating	4.53 (0.58)	4.53 (0.56)	4.55 (0.60)	4.50 (0.63)	4.53 (0.59)
Located in IDF	0.99 (0.10)	0.97 (0.16)	1.00 (0.07)	0.92 (0.27)	0.97 (0.18)
Located in Paris	0.74 (0.44)	0.78 (0.42)	0.90 (0.30)	0.80 (0.40)	0.81 (0.39)
Crow distance	17.31 (44.95)	17.85 (53.39)	15.39 (47.70)	17.24 (51.94)	17.21 (51.95)
Same department	0.29 (0.45)	0.34 (0.48)	0.33 (0.47)	0.33 (0.47)	0.34 (0.47)
Same postcode	0.01 (0.12)	0.02 (0.14)	0.02 (0.13)	0.02 (0.13)	0.02 (0.13)
Time since last connection					
Connected in the last hour	0.56	0.58	0.51	0.50	0.54
Connected in the 6 hours	0.21	0.22	0.22	0.23	0.22
Connected in the 12 hours	0.09	0.08	0.09	0.09	0.09
Connected in the last day	0.08	0.07	0.10	0.09	0.08
Connected in last week	0.05	0.04	0.07	0.07	0.06
No recent connection	0.01	0.01	0.01	0.01	0.01
N distinct individuals	4,007	2,307	2,519	1,716	4,539
N	174,577	79,894	28,895	37,467	146,256
Panel B: firm characteristics					
Nb. of venues	1.52 (2.44)	1.21 (1.18)	1.22 (1.04)	1.19 (0.89)	1.20 (1.05)
Nb. of postings	11.51 (28.54)	13.94 (30.13)	10.57 (19.45)	10.22 (21.57)	12.04 (25.85)
N distinct establishments	1,101	540	141	450	848
N	12,677	7,530	1,491	4,599	13,620

Notes: The table reports applicant and employer characteristics for waiting-staff, bar-staff, receptionist, and runner positions advertised between 5 April 2023 and 4 May 2024. Standard deviations are reported in parentheses below means. No additional statistic is reported below proportions.

Figure 1.3 plots monthly median and mean hourly wages in the four occupations. Because most employers use the platform default, the median generally equals the default wage. The default remained unchanged throughout the sample period for receptionists, runners, and bar staff. On 5 October 2023, however, the waiting-staff default rose from €16 to €17, an increase of 6.25%. The platform notified client firms of the change.

Mean wages are noisier because employers can use client-specific wage grids or boost individual postings. Boosting is particularly visible for bar-staff and re-


1.2. INSTITUTIONAL CONTEXT AND DATA

21

ceptionist jobs in December 2023. Nevertheless, both panels show a clear increase in waiting-staff wages relative to wages in the three comparison occupations.

We use this relative wage change to estimate effects on the quantity and quality of applications and on hiring.

Figure 1.3: Hourly wages by month and occupation



Notes: The figure reports median and mean advertised hourly wages for waiting-staff, bar-staff, receptionist, and runner jobs between 5 April 2023 and 4 May 2024. Values are calculated from the job-level data.

1.3 Conceptual framework

We develop a model of application and hiring decisions on the platform. The model serves two purposes. First, it clarifies why a higher wage need not increase applications from every type of worker. A worker applies to a job not only because it pays well, but also because the worker has a reasonable chance of obtaining it. A wage increase can attract additional high-productivity applicants, lower the hiring prospects of other applicants, and discourage some of them from applying. Second, the model shows how a wage change in one occupation affects applications to other occupations through the value of continuing to search.

We abstract from several institutional details. The framework contains two occupations: occupation A , which represents waiting-staff jobs and whose wage changes in the empirical analysis, and occupation B , which represents the pooled comparison occupations. In Section 1.6, we estimate the model's parameters and use them to decompose the reduced-form estimates into their equilibrium components. Appendix 1.C provides more detail about the conceptual framework, as well as the proofs.

1.3.1 Environment

Time is discrete. Workers are risk neutral and discount future utility by $\tilde{\beta}$. We focus on workers who are unemployed and available for work on the platform. Such a worker has type

$$q = (q_A, q_B) \in [0, 1]^2,$$

where q_j is the worker's productivity rank in occupation j . Higher values indicate higher productivity. Let F denote the stationary joint distribution of types among workers who are unemployed and available on the platform. The joint distribution allows workers who are productive in one occupation to be systematically more or less productive in the other. We take this stationary distribution and the mass of unemployed available workers as primitives and hold them fixed in the wage comparative statics; the framework therefore abstracts from changes in unemployment stocks and their composition.

The timing within a period is as follows. An unemployed worker first receives the flow payoff b . The worker then receives at most one notification: a notification from occupation $j \in \{A, B\}$ arrives with type-specific probability $\alpha_j(q)$, where $\alpha_A(q) + \alpha_B(q) \leq 1$. Upon receiving a notification, the worker observes a match-specific payoff shock ε_j , pays $\tilde{c}_j$ if applying, and learns the hiring outcome before the end of the period. A successful applicant begins employment in the next period. The shock ε_j is logistic with scale s_j , is independent of the employer's

1.3. CONCEPTUAL FRAMEWORK

23

evaluation shock, and remains fixed over the employment spell. Employment yields $w_j + \varepsilon_j$ per period and ends at the close of each period with probability δ , after which the worker returns to unemployment and becomes available on the platform again.

For each occupation- j vacancy, the number of notified workers is Poisson with mean M_j . Thus M_j is a notification intensity, rather than a fixed number of workers. Let

$$\bar{\alpha}_j = \int_{[0,1]^2} \alpha_j(q) dF(q)$$

be the average probability that an unemployed available worker receives a j -notification.

If N^U is the mass of unemployed available workers and V_j is the mass of occupation- j vacancies in the same period, the two ways of counting notifications imply the market identity

$$V_j M_j = N^U \bar{\alpha}_j, \quad \text{or equivalently} \quad M_j = \frac{N^U}{V_j} \bar{\alpha}_j. \quad (1.1)$$

Conditional on an occupation- j notification, the worker's type is distributed according to

$$dF_j^N(q) = \frac{\alpha_j(q)}{\bar{\alpha}_j} dF(q). \quad (1.2)$$

Employers value both observed productivity and an idiosyncratic employer-applicant match component. Specifically, employer j 's payoff from applicant i is $q_{ij} + v_{ij}$, where v_{ij} follows a type-I extreme-value distribution with scale σ_j . The employer selects the applicant with the highest payoff, and the selected applicant is successfully matched with probability κ_j . Thus σ_j governs how strongly employers sort applicants by observed productivity: selection approaches deterministic ranking when σ_j is small and random selection when it is large. We derive the resulting selection probabilities below.

1.3.2 Application decisions

Let $U(q; w)$ be the value of remaining unemployed and available at the beginning of a period, where $w = (w_A, w_B)$, and define the flow-equivalent continuation value

$$u(q; w) = (1 - \tilde{\beta})U(q; w).$$

The flow value of remaining available for work is the constant b . The structural specification in Section 1.6 allows this value to vary across worker types. Let $p_j(q; w)$ denote the equilibrium probability that a worker is hired conditional on applying to occupation j . After rescaling the application cost by the expected

duration of employment, the worker applies if and only if

$$p_j(q_j; w) [w_j + \varepsilon_j - u(q; w)] \geq c_j. \quad (1.3)$$

Conditional on a notification from occupation j , a worker's application probability is

$$\lambda_j(q; w) = \frac{1}{1 + \exp \{ -[w_j - u(q; w) - c_j / p_j(q_j; w)] / s_j \}}. \quad (1.4)$$

Under equation (1.4), the application behaviour is different from a standard reservation-wage model. With a fixed hiring probability, the worker compares the wage with a reservation value. Here that reservation value is itself higher when the job is harder to obtain:

$$\text{effective reservation value} = u(q; w) + \frac{c_j}{p_j(q_j; w)}.$$

A low-productivity worker may therefore decline a high-wage opportunity if the expected cost of securing the job is sufficiently large.

The continuation value links the two occupations. Under the timing above, with logistic shocks it can be written exactly as

$$u(q; w) = b - \sum_{k \in \{A, B\}} \alpha_k(q) \beta s_k p_k(q_k; w) \log[1 - \lambda_k(q; w)], \quad \beta = \frac{\tilde{\beta}}{1 - \tilde{\beta}(1 - \delta)}. \quad (1.5)$$

Here β is the discounted value of a unit flow received during an employment spell that begins next period; it is not a second primitive discount factor. Appendix 1.C derives this expression and the logistic option-value term. Better opportunities in A raise the value of waiting for an A notification. This makes a current opportunity in B less attractive, even though w_B has not changed.

1.3.3 Equilibrium competition

Application decisions determine the applicant pool and hence the probability of being hired. Write $a_j(q_j) = \exp(q_j / \sigma_j)$, and let $S_j(w)$ denote the random sum of $a_j(Q_j)$ across the other applicants to a vacancy. Because the employer's match component is type-I extreme-value, the exact conditional hiring probability of a focal applicant is

$$p_j^E(q_j; w) = \kappa_j \mathbb{E} \left[\frac{a_j(q_j)}{a_j(q_j) + S_j(w)} \middle| q_j, \text{ the worker applies} \right]. \quad (1.6)$$

1.3. CONCEPTUAL FRAMEWORK

25

Appendix 1.C gives the exact Poisson-logit operator. To obtain a transparent expression in the main text, we use a mean-field approximation that replaces $S_j(w)$ by its expectation. This gives

$$p_j^{MF}(q_j; w) = \kappa_j \frac{a_j(q_j)}{a_j(q_j) + D_j(w)}, \quad (1.7)$$

$$\begin{aligned} D_j(w) &= M_j \int_{[0,1]^2} a_j(q_j) \lambda_j(q; w) dF_j^N(q) \\ &= \frac{M_j}{\bar{\alpha}_j} \int_{[0,1]^2} \alpha_j(q) a_j(q_j) \lambda_j(q; w) dF(q). \end{aligned} \quad (1.8)$$

The index $D_j(w) = \mathbb{E}[S_j(w)]$ is expected effective competition from other applicants. Since $s \mapsto a/(a+s)$ is convex for $a > 0$, Jensen's inequality implies

$$p_j^E(q_j; w) \geq p_j^{MF}(q_j; w), \quad (1.9)$$

with equality only when the random competitor index $S_j(w)$ is degenerate. Thus the mean-field expression is a lower bound on the exact hiring probability, not an exact ratio of expectations. In the main-text approximation, the p_j entering equations (1.3)–(1.5) is set equal to p_j^{MF} . A mean-field equilibrium consists of application probabilities, mean-field hiring probabilities, and continuation values satisfying equations (1.4)–(1.8). Appendix 1.C defines the exact equilibrium, derives the Jensen gap to second order, and distinguishes exact from mean-field limiting cases. The wage response of applications combines three forces.

Proposition 1.1 (Wages, competition, and application decisions). *For a worker of type q ,*


$$\frac{d\lambda_j(q; w)}{dw_k} = \frac{\lambda_j(q; w)[1 - \lambda_j(q; w)]}{s_j} \left[\mathbf{1}\{j = k\} - \frac{du(q; w)}{dw_k} + \frac{c_j}{p_j(q_j; w)^2} \frac{dp_j(q_j; w)}{dw_k} \right]. \quad (1.10)$$

Consequently, an increase in w_j need not increase applications to occupation j for every worker type.

For the occupation whose wage rises, the first term in brackets is the direct wage effect. The second term is an option-value effect: a higher wage raises the value of remaining available for future opportunities and therefore makes the current job less attractive. The third term is the competition effect. If the wage attracts additional applicants, D_j rises and p_j falls. The resulting increase in c_j/p_j is especially important for workers with low hiring probabilities. For these workers, the competition and option-value effects can outweigh the direct wage effect.

Figure 1.4 illustrates this possibility. Panel (a) compares application profiles

Figure 1.4: Wages, hiring prospects, and application decisions



Notes: The figure solves the model under an illustrative calibration. Panel (a) reports the average probability of applying to occupation A, conditional on an A notification, at two wages. Panel (b) reports the same probability as a function of w_A for three productivity ranks. Continuation values and hiring probabilities adjust in equilibrium along each curve.

at two wages. The profiles cross because the wage increase raises applications among high-productivity workers but lowers them for some lower-productivity workers. Panel (b) follows three worker types over a wider wage range. Applications rise monotonically for the high-productivity type. For the two lower-productivity types, applications first rise with the wage and then fall once the induced increase in competition dominates the direct wage incentive. The numerical values are chosen only to display the mechanism and are not estimates.

The same expression clarifies spillovers. When w_A rises, applications to B have no direct wage term:

$$\frac{d\lambda_B}{dw_A} = \frac{\lambda_B(1 - \lambda_B)}{s_B} \left[-\frac{du}{dw_A} + \frac{c_B}{p_B^2} \frac{dp_B}{dw_A} \right].$$

The increase in the continuation value tends to reduce applications to B. As those applications fall, competition for B jobs weakens and p_B rises, which mitigates the initial decline. The size and composition of the spillover depend on the joint distribution of q_A and q_B : it is strongest when workers who are marginal applicants to B also have good prospects in A.

1.3. CONCEPTUAL FRAMEWORK

27

1.3.4 Observable outcomes

The model maps application decisions into the three classes of outcomes studied empirically. Expected applications per vacancy are

$$\mathcal{A}_j(w) = M_j \int_{[0,1]^2} \lambda_j(q;w) dF_j^N(q) = \frac{M_j}{\bar{\alpha}_j} \int_{[0,1]^2} \alpha_j(q) \lambda_j(q;w) dF(q). \quad (1.11)$$

The productivity distribution among applicants is the notified type distribution reweighted by $\lambda_j(q;w)$. In particular, mean applicant quality is

$$\bar{Q}_j(w) = \frac{\int_{[0,1]^2} q_j \lambda_j(q;w) dF_j^N(q)}{\int_{[0,1]^2} \lambda_j(q;w) dF_j^N(q)}. \quad (1.12)$$

The appendix gives the corresponding expressions for other functionals of the applicant pool, including its minimum and maximum quality.

Finally, under the exact Poisson model, application arrivals imply that the job-filling rate is


$$H_j(w) = \kappa_j [1 - \exp\{-\mathcal{A}_j(w)\}]. \quad (1.13)$$

Equivalently, under the exact hiring operator, the application-weighted individual hiring probabilities add up to $H_j(w)$. The mean-field lower bound is used to approximate individual hiring prospects and does not in general preserve this exact adding-up identity; Appendix 1.C makes the distinction explicit. Applications therefore raise filling mechanically, but the marginal effect becomes small when vacancies already receive many applications. Applicant quality has no similarly mechanical response: it rises when the wage draws disproportionately more high-productivity workers or discourages low-productivity workers, and falls under the opposite compositional change.

Figure 1.5 illustrates these mappings. In the example, a higher w_A strongly raises applications and average applicant quality in A , while slightly reducing applications to B .

These equilibrium links matter for the difference-in-differences design. The observed change for occupation A includes the direct wage incentive and the endogenous responses of competition and continuation values. Moreover, occupation B need not be unaffected: the improved opportunity in A changes workers' outside option and can alter applications, applicant quality, and filling in B . The structural analysis uses the estimated model to separate these components.

Figure 1.5: From wages to applications and applicant quality



Notes: The figure uses the same illustrative calibration as Figure 1.4. Occupation B’s wage is held fixed. Applicant quality is the mean occupation-specific productivity rank among applicants.

1.4. THE IMPACT OF WAGES ON APPLICATIONS

29

1.4 The impact of wages on applications

This section estimates the effect of the waiting-staff wage increase on the quantity and quality of applications, using the other three occupations as the comparison group.

1.4.1 Empirical strategy

We compare outcomes before and after 5 October 2023 for waiting-staff jobs and jobs in the three comparison occupations. Because employers can depart from the platform default, the estimates are intention-to-treat effects.¹

For the number of applications A_j received by job j , we estimate the following Poisson model:

$$\mathbb{E}[A_j | X_j] = \exp\left(\beta Z_{o(j)} \mathbb{1}\{t(j) > t_0\} + X_j' \alpha + \lambda_{o(j)} + \lambda_{\text{week}(j)} + \phi_j + \ln M_j\right). \quad (1.14)$$

For an applicant-quality outcome Y_j , we estimate the corresponding linear model:

$$Y_j = \beta Z_{o(j)} \mathbb{1}\{t(j) > t_0\} + X_j' \alpha + \lambda_{o(j)} + \lambda_{\text{week}(j)} + \phi_j + \varepsilon_j. \quad (1.15)$$

Here $o(j)$ and $t(j)$ denote the occupation and creation date of job j , $Z_{o(j)}$ equals one for waiting-staff jobs, and t_0 is the date of the wage change. The vector X_j includes indicators for morning, lunch, afternoon, evening, and night shifts; contracted work and break hours; the day of the month, day of the week, and hour at which the advert was created; deciles of the time between posting and the start of the job; indicators for the number of working days spanned by the job; and indicators for weekday-only, weekend-only, and mixed jobs. In equation (1.14), $\ln M_j$ is an offset for the number of workers notified about job j .

All specifications include occupation and week fixed effects. The additional fixed effects in ϕ_j vary across columns. Our preferred specification includes venue fixed effects, where a venue is an establishment–postcode pair. The coefficient β captures the intention-to-treat effect of the default-wage increase. For quality, we consider the mean, minimum, and maximum score among applicants to each job.

¹In the analysis sample, 95% of adverts post the occupation-wide default wage, 2% post a lower wage, and 3% post a higher wage.

We also estimate event-study versions of equations (1.14) and (1.15):

$$\mathbb{E}[A_j | X_j] = \exp \left(\sum_{k=-6, k \neq -1}^6 \beta_k Z_{o(j)} \mathbb{1}\{m(j) = k\} + X_j' \alpha + \lambda_{o(j)} + \lambda_{\text{week}(j)} + \phi_j + \ln M_j \right) \quad (1.16)$$

$$Y_j = \sum_{k=-6, k \neq -1}^6 \beta_k Z_{o(j)} \mathbb{1}\{m(j) = k\} + X_j' \alpha + \lambda_{o(j)} + \lambda_{\text{week}(j)} + \phi_j + \varepsilon_j \quad (1.17)$$

where $m(j)$ indexes event months relative to t_0 ; the month from 5 October to 4 November 2023 is zero, the next month is one, and so forth. The omitted category is the month immediately preceding the change. Defining months around the event date ensures that the default wage is constant within each event month.

Identification requires that, absent the wage change, outcomes in the treated and comparison occupations would have followed parallel trends. Occupation-specific trends or seasonal patterns could violate this condition. Detailed job controls and venue fixed effects reduce concerns about changes in the composition of postings. Moreover, many of the same firms post jobs, and many of the same workers apply, across the four occupations. Common demand- or supply-side shocks are therefore likely to affect more than one occupation.

This overlap also creates interference: additional applications to waiting-staff jobs may come at the expense of the comparison occupations. The reduced-form coefficient consequently identifies a relative effect, not necessarily the equilibrium response in the treated occupation alone. Section 1.6 uses the model to separate these objects.

1.4.2 Defining quality

We construct quality from employers' revealed preferences over applicants. Specifically, we use a class-weighted logistic model to predict whether worker i is hired for job j , conditional on applying. Because the estimation reweights hires and non-hires to address class imbalance, the fitted value is best interpreted as a hiring score rather than as a calibrated probability. Quality is worker-job specific: the same worker can have different scores for different jobs.

$$S_{ij} = \frac{1}{1 + \exp(-z_{ij})}, \quad z_{ij} = \beta_0 + X_{ij}' \beta + \gamma_{c(j),e(j)}, \quad (1.18)$$

1.4. THE IMPACT OF WAGES ON APPLICATIONS


31

where S_{ij} is the quality score, X_{ij} contains worker, job, and worker–job characteristics, and $\gamma_{c(j),e(j)}$ denotes establishment fixed effects indexed by establishment category and employer. The predictors include whether the worker has previously worked for and been approved by the employer, distance between the worker's residence and the workplace, the worker's average rating, experience in the job's occupation and establishment category, gender, age, and time since the worker last connected to the platform. We estimate the model using service-sector applications between 1 January 2023 and 30 December 2025. Appendix 1.B describes the sample, predictors, and estimation in detail.

We draw 20% of jobs at random to form the training sample and calculate scores for applications to the remaining 80%. The quality regressions in equations (1.15) and (1.17) exclude all jobs in the training sample, preventing a job's realised hiring outcome from being used both to estimate and to evaluate its applicants' quality scores.

In the test sample, the area under the receiver operating characteristic curve (AUC) is 0.76. At the selected threshold, sensitivity—the share of hires classified as hires—is 0.61. Figure 1.6 plots the score distribution separately for hired and non-hired applicants. Although the distributions overlap, hired applicants receive substantially higher scores.

Figure 1.6: Applicant quality scores by hiring outcome



Notes: The figure reports the density of applicant quality scores by hiring outcome in the test sample. The sample contains applications to service-sector jobs posted between 1 January 2023 and 30 December 2025, excluding jobs used to train the quality model.

1.4.3 Results

Table 1.3 reports estimates of equation (1.14). Column (1) includes occupation and week fixed effects; column (2) adds establishment-category and postcode fixed ef-

fects; column (3) adds their interaction; and column (4) instead includes venue fixed effects. Across specifications, the wage increase raised applications by 22–30% relative to the comparison occupations.² In our preferred specification, $\hat{\beta} = 0.2104$, corresponding to a 23.4% increase. The log elasticity is $0.2104 / \log(17/16) = 3.5$.

Table 1.3: Effect of the wage increase on applications

	Nb. of Applications			
	(1)	(2)	(3)	(4)
Post × Treated	0.2662*** (0.0204)	0.1950*** (0.0200)	0.2066*** (0.0209)	0.2104*** (0.0221)
Observations (# of Missions)	26,297	26,112	25,921	25,461
Share of Treated Observations	0.482	0.48	0.478	0.474
Dependent Variable Mean	12.2	12.2	12.3	12.2
Number of Venues	2,200	2,021	1,832	1,379
Controls Included	✓	✓	✓	✓
Occupation FE	✓	✓	✓	✓
Week FE	✓	✓	✓	✓
Category FE		✓	✓	
Postcode FE		✓	✓	
Category × Postcode FE			✓	
Venue FE				✓


Figure 1.7 presents the event-study estimates from the preferred specification of equation (1.16). The pre-treatment coefficients are small and statistically insignificant, providing no evidence of differential application trends before the wage change. Applications rise gradually after the change, reaching approximately 20% above the relative pre-treatment level after two months. The effect persists through the end of the six-month post-treatment period.

²Percentage changes are calculated as $100[\exp(\hat{\beta}) - 1]$.

1.4. THE IMPACT OF WAGES ON APPLICATIONS

33

Figure 1.7: Event study: number of applications



Notes: The figure reports the estimated β_k coefficients from equation (1.16). The model includes occupation, week, and venue fixed effects and uses the number of notified workers as an offset. The sample runs from 5 April 2023 to 4 May 2024. Bars show 95% confidence intervals.

Table 1.4 reports effects on the mean, minimum, and maximum quality scores in the applicant pool. Mean quality increases by 0.007–0.013 across specifications. In the preferred specification, the estimate is 0.0095, or 2.0% of the dependent-variable mean, implying an elasticity of approximately 0.3. The minimum score rises by 0.0307, or 10.6% of its mean, implying an elasticity of approximately 1.7. By contrast, the maximum-score estimate is 0.0065, or 0.9% of its mean, and is statistically insignificant. The improvement in applicant quality is therefore concentrated at the lower end of the distribution.


CHAPTER 1. THE WAGE ELASTICITY OF JOB APPLICATIONS

Table 1.4: Effect of the wage increase on applicant quality


	(1)	(2)	(3)	(4)
<i>Panel A. Mean quality</i>				
Post × Treated	0.0084*** (0.0031)	0.0145*** (0.0031)	0.0122*** (0.0032)	0.0109*** (0.0035)
Dependent Variable Mean	0.479	0.479	0.480	0.480
<i>Panel B. Minimum quality</i>				
Post × Treated	0.0279*** (0.0037)	0.0334*** (0.0038)	0.0312*** (0.0040)	0.0283*** (0.0044)
Dependent Variable Mean	0.276	0.276	0.276	0.276
<i>Panel C. Maximum quality</i>				
Post × Treated	0.0067 (0.0053)	0.0114** (0.0051)	0.0078 (0.0053)	0.0086 (0.0057)
Dependent Variable Mean	0.727	0.727	0.728	0.729
Observations (# of Missions)	19,503	19,388	19,207	18,773
Share of Treated Observations	0.491	0.490	0.488	0.483
Number of Venues	1,915	1,800	1,619	1,185
Controls Included	✓	✓	✓	✓
Occupation FE	✓	✓	✓	✓
Week FE	✓	✓	✓	✓
Category FE		✓	✓	
Postcode FE		✓	✓	
Category × Postcode FE			✓	
Venue FE				✓

Figure 1.8 presents the event-study counterparts of column (4) of each panel of Table 1.4. The increase in mean quality is short-lived and statistically significant only in the first two months after the event month. Minimum quality rises more sharply and remains significantly higher throughout the post-treatment period. Maximum quality shows no clear or persistent response.


Figure 1.8: Event study: applicant quality



(a) Mean applicant quality



(b) Minimum applicant quality



(c) Maximum applicant quality

Notes: Panels 1.8a–1.8c report the estimated β_k coefficients from equation (1.17) for mean, minimum, and maximum applicant quality. All specifications include occupation, week, and venue fixed effects. The sample runs from 5 April 2023 to 4 May 2024. Bars show 95% confidence intervals.

1.5 The impact of wages on hires

We apply the same difference-in-differences and event-study designs to the number of workers hired for each job. Because the default wage is not binding, these estimates are again intention-to-treat effects.

Identification requires parallel trends in hiring absent the wage change. The threats are similar to those discussed in Section 1.4, and we use the same job controls and fixed effects. Hiring differs from applying, however, because it reflects the joint decisions of workers and employers. Employers may respond to the wage increase by changing their selection standards or recruitment practices. The estimated effect on hires therefore combines the application response with any wage-induced change in employers' hiring decisions.

Table 1.5 reports the effect on the number of hires. Across specifications, the wage increase raises hires by 11–16% relative to the comparison occupations. In the preferred specification, $\hat{\beta} = 0.1087$, corresponding to an 11.5% increase. The implied log elasticity is 1.8. Figure 1.9 presents the corresponding event study. The post-treatment coefficients are generally positive but individually imprecise, so none is statistically distinguishable from zero. The pooled post-treatment estimate in Table 1.5, by contrast, is positive and statistically significant.


Table 1.5: Effect of the wage increase on hires

	Nb. of Hires			
	(1)	(2)	(3)	(4)
Post $\times$ Treated	0.1028*** (0.0338)	0.1086*** (0.0342)	0.1481*** (0.0354)	0.1087*** (0.0375)
Observations (# of Missions)	26,297	25,938	25,509	24,629
Share of Treated Observations	0.482	0.478	0.476	0.471
Dependent Variable Mean	0.6	0.6	0.6	0.6
Number of Venues	2,200	1,952	1,694	1,137
Controls Included	✓	✓	✓	✓
Occupation FE	✓	✓	✓	✓
Week FE	✓	✓	✓	✓
Category FE		✓	✓	
Postcode FE		✓	✓	
Category $\times$ Postcode FE			✓	
Venue FE				✓

1.5. THE IMPACT OF WAGES ON HIRES

37

Figure 1.9: Event study: number of hires



Notes: The figure reports the estimated β_k coefficients from equation (1.16) for the number of hires. The model includes occupation, week, and venue fixed effects and uses the number of notified workers as an offset. The sample runs from 5 April 2023 to 4 May 2024. Bars show 95% confidence intervals.

1.6 Structural decomposition of the reduced-form estimators

We estimate the model developed in Section 1.3. Our main purpose is to decompose the difference-in-differences estimates into the equilibrium effect on the occupation in which the wage increased and the displacement effect on the comparison occupations. A secondary benefit is that the richness of the platform data identifies structural parameters that are typically difficult to recover, including application costs and the dispersion in workers' and employers' evaluations of a match. Details about the estimation and the decomposition are available in Appendix 1.D.

1.6.1 Parameters and moments

The structural model discretises the stationary distribution F into cells c of the joint (q_A, q_B) distribution. Let ω_c be the share of unemployed workers available on the platform in cell c , and let $\alpha_j(c)$ be their per-period probability of an occupation- j notification. The average notification probability and the type distribution conditional on a j -notification are

$$\bar{\alpha}_j = \sum_c \omega_c \alpha_j(c), \quad \pi_j(c) = \frac{\omega_c \alpha_j(c)}{\bar{\alpha}_j}. \quad (1.19)$$

Consequently, the structural analogues of applications and effective competition are

$$A_j = M_j \sum_c \pi_j(c) \lambda_j(c), \quad D_j = M_j \sum_c \pi_j(c) \exp\{q_j(c)/\sigma_j\} \lambda_j(c). \quad (1.20)$$

This is a numerical discretisation of the integrals in Section 1.3, rather than an additional economic assumption.

The worker-side and vacancy-side notification objects are not independent free parameters. If N^U is the number of unemployed available worker-periods and V_j is the number of occupation- j vacancies in the same sampling window, they must satisfy

$$V_j M_j = N^U \sum_c \omega_c \alpha_j(c). \quad (1.21)$$

We therefore construct M_j , $\alpha_j(c)$, and ω_c from a common notification universe and verify equation (1.21) before estimation. Equivalently, one side of the identity may be computed from the other. An unreconciled pair of $\alpha_j(c)$ and M_j is not an admissible calibration.

The model contains three groups of parameters. The first group is observed or

1.6. STRUCTURAL DECOMPOSITION

39

calibrated directly:

$$\theta_0 = (M_A, M_B, \{\alpha_A(c), \alpha_B(c)\}_c, \delta, \tilde{\beta}, \{\omega_c, q_A(c), q_B(c), \chi_A(c), \chi_B(c)\}_c).$$

Here M_j is the mean Poisson notification intensity per vacancy, $\chi_j(c)$ is the mean empirical hiring score associated with cell c in occupation j , and the remaining objects have their definitions above. Mean match duration is observed on the platform, and the discount factor is fixed from an annual discount rate. Earlier versions of the quantitative model used a Gaussian copula for the joint distribution of ranks; the baseline instead uses the empirical joint cells directly.

The second group contains occupation-specific matching efficiencies (κ_A, κ_B) . Let $\mathcal{A}_j$ denote applications per vacancy and z_j the fraction of vacancies that remain unfilled. Because application arrivals are Poisson,

$$\kappa_j = \frac{1 - z_j}{1 - \exp(-\mathcal{A}_j)}. \quad (1.22)$$

We recover κ_j from pre-reform applications and filling rates. The model treats each job as one effective vacancy and defines filling as at least one hire. It therefore abstracts from jobs that request or hire several workers. When relating the model to the reduced-form estimate for the number of hires, we assume that the elasticity of hires is well approximated by the elasticity of the probability of at least one hire. This restriction is transparent but is not used to reinterpret a multi-worker job as several independent vacancies.

The remaining behavioural parameters are estimated jointly:

$$\theta = (b, b_1, c_A, c_B, s_A, s_B, \sigma_A, \sigma_B)'. \quad (1.23)$$

For the structural specification, we generalise the constant unemployment flow payoff in Section 1.3 to

$$b(c) = b + b_1 \bar{q}(c), \quad \bar{q}(c) = \frac{q_A(c) + q_B(c)}{2}. \quad (1.24)$$

The intercept b governs the attractiveness of remaining available for future work, and b_1 allows this value to vary with average productivity rank. The c_j 's are application costs, the s_j 's govern dispersion in workers' application payoffs, and the σ_j 's govern dispersion in employers' evaluations of applicants.

We construct pre-reform moments by crossing bins of q_A and q_B . For every joint worker cell c , we measure the probability of applying conditional on receiving a

notification:

$$\hat{\lambda}_j(c) = \frac{\text{applications from cell } c \text{ to } j}{\text{notifications received by cell } c \text{ from } j}.$$

We also measure hiring probabilities conditional on applying. Consistent with the model, p_A is estimated by q_A bin and p_B by q_B bin:

$$\hat{p}_j(x) = \frac{\text{hires from applications in rank bin } x}{\text{applications in rank bin } x}.$$

The pre-reform moments additionally include applications per vacancy in each occupation. These aggregate moments are computed directly from job-level application and vacancy totals in a held-out aggregation fold; they are not reconstructed from M_j and the cell application moments. They therefore provide genuine overidentifying information rather than duplicating an accounting identity. Appendix 1.D describes the sample, weights, standard errors, and construction of the representative comparison occupation.

Finally, we add the empirical difference-in-differences estimates for applications and the average applicant hiring score. Productivity ranks $q_j(c)$ enter worker heterogeneity and employer sorting, but the empirical quality outcome is the score S_{ij} , not its percentile rank. The structural counterpart of the empirical mean score is therefore

$$\bar{S}_j(w) = \frac{\sum_c \pi_j(c) \chi_j(c) \lambda_j(c; w)}{\sum_c \pi_j(c) \lambda_j(c; w)}. \quad (1.25)$$

Let w_A^0 be the pre-reform wage and $w_A^1 = w_A^0 + \Delta w_A$ the post-reform wage. Solving the model at both wages gives

$$\tau_A^M(\theta) = [h(\mathcal{A}_A^1) - h(\mathcal{A}_A^0)] - [h(\mathcal{A}_B^1) - h(\mathcal{A}_B^0)], \quad (1.26)$$

$$\tau_S^M(\theta) = [\bar{S}_A^1 - \bar{S}_A^0] - [\bar{S}_B^1 - \bar{S}_B^0], \quad (1.27)$$

where h applies the same transformation as the empirical specification. Because the application regression is Poisson, its coefficient is a change in the log application rate, so its structural counterpart uses $h(x) = \log x$. The application equation and employer-choice index continue to use percentile ranks, whereas the quality DiD is matched in the original score units through equation (1.25). The distinction is important because changes in means are not invariant to a rank transformation.

The baseline estimator uses the mean-field hiring probabilities in equation (1.7). As a specification exercise, we also re-estimate the model after replacing the mean-field update with the exact Poisson-logit operator in equation (1.43). All empirical moments and other primitives are held fixed across the two estimators. The comparison reports changes in fit, parameters, hiring probabilities, and the decompo-

1.6. STRUCTURAL DECOMPOSITION

41

sition; it is a sensitivity analysis, not a mixture of the two operators within one equilibrium.

1.6.2 Identification

Identification follows the sequence suggested by the economic structure. Match duration identifies δ , an externally specified annual discount rate calibrates $\tilde{\beta}$, and notification counts identify $\alpha_j(c)$, M_j , and the notified-worker shares subject to the reconciliation restriction (1.21). Filling and application rates identify κ_j through equation (1.22), and the empirical cell weights identify the joint distribution of productivity ranks used in the equilibrium computation.

The slope of hiring probabilities across productivity ranks identifies employer selectivity. Rearranging the mean-field hiring equation yields

$$\log \left[\frac{\kappa_j}{p_j(q_j)} - 1 \right] = \log D_j - \frac{q_j}{\sigma_j}. \quad (1.28)$$

Thus the slope of transformed hiring probabilities in q_j is $-1/\sigma_j$. A flatter hiring profile implies more employer-side noise and a larger σ_j .

Conditional application probabilities identify the worker-side parameters. Define

$$G_j(q) = -\alpha_j(q)\beta p_j(q_j) \log[1 - \lambda_j(q)].$$

Combining the application rule with the continuation-value equation gives

$$w_j = b + b_1 \bar{q} + \frac{c_j}{p_j(q_j)} + s_j \text{logit}[\lambda_j(q)] + s_A G_A(q) + s_B G_B(q). \quad (1.29)$$

Across occupations and productivity cells, equation (1.29) is linear in $(b, b_1, c_A, c_B, s_A, s_B)$. These parameters are point identified when the resulting regressor matrix has rank six. Intuitively, variation in hiring probabilities identifies application costs, variation in application rates identifies the worker-side shock scales, and behaviour across both occupations and quality cells identifies the level and type gradient of the continuation value.

The pre-reform conditional moments are sufficient for point identification under these rank conditions. The two difference-in-differences moments add economically important overidentifying information. They discipline the parameter combinations that determine continuation-value and competition responses to wages and require the model's comparative statics to align with the reduced-form evidence. Appendix 1.D states the formal identification result.

1.6.3 GMM estimator

Let $\hat{m}$ collect the pre-reform cell and aggregate moments and the two difference-in-differences estimates, and let $m(\theta)$ denote their model counterparts. We standardise each discrepancy by its estimated standard error and stack the resulting moments in $g(\theta)$. The estimator is

$$\hat{\theta} = \arg \min_{\theta \in \Theta} g(\theta)' W g(\theta). \quad (1.30)$$

We impose positivity of costs and scale parameters by optimising their logarithms. The intended estimator uses the inverse of the jointly estimated covariance matrix of all moments. A common resampling scheme recomputes the calibrated primitives, pre-reform moments, and difference-in-differences estimates in every draw, preserving their covariance. Thus uncertainty from the notification reconciliation, generated ranks and scores, and calibrated κ_j enters both estimation and inference.

With this joint covariance matrix, the overidentifying-restrictions statistic is

$$J = g(\hat{\theta})' W g(\hat{\theta}),$$

which is asymptotically chi-squared with the number of moments minus eight degrees of freedom under the maintained model and standard regularity conditions. We report both the asymptotic p -value and a bootstrap p -value obtained by re-estimating the model in each draw. A diagonal or partially supplied covariance matrix is useful only as a sensitivity exercise and does not deliver the baseline J -test. The same common bootstrap provides sampling uncertainty for the parameters and decomposition.

1.6.4 Decomposition

For outcome Y , let $Y_j(w_A; \hat{\theta})$ be the model outcome in occupation j when the wage in A is w_A . Holding w_B , the worker distribution, notification rates, and all estimated parameters fixed, define

$$\Delta Y_j = Y_j(w_A^1; \hat{\theta}) - Y_j(w_A^0; \hat{\theta}).$$

The structural counterpart of the empirical difference-in-differences coefficient is

$$\tau_Y^M = \underbrace{\Delta Y_A}_{\text{effect on treated occupation}} + \underbrace{(-\Delta Y_B)}_{\text{control-group spillover}}. \quad (1.31)$$

1.6. STRUCTURAL DECOMPOSITION

43

The first term is the full equilibrium response in A . It includes the direct wage incentive as well as endogenous changes in competition and continuation values. The second term measures contamination of the comparison group. If workers substitute away from B , then $\Delta Y_B < 0$ for applications and the DiD exceeds the equilibrium effect on A .

We further decompose the spillover on B . Along the equilibrium wage path,

$$\frac{d\lambda_B(q)}{dw_A} = \underbrace{-\frac{\lambda_B(q)[1 - \lambda_B(q)]}{s_B} \frac{du(q)}{dw_A}}_{\text{continuation-value channel}} + \underbrace{\frac{\lambda_B(q)[1 - \lambda_B(q)]}{s_B} \frac{c_B}{p_B(q_B)^2} \frac{dp_B(q_B)}{dw_A}}_{\text{hiring-probability channel}}. \quad (1.32)$$

The first channel normally reduces applications to B : a better opportunity in A raises the value of continuing to search. The second channel often mitigates that reduction: fewer applications to B lower competition and raise the probability of being hired there.

We aggregate each derivative over worker types and integrate it from w_A^0 to w_A^1 . For applications and average quality this yields

$$\Delta Y_B = \Delta Y_B^u + \Delta Y_B^p, \quad \tau_Y^M = \Delta Y_A - \Delta Y_B^u - \Delta Y_B^p.$$

The decomposition is performed in the same units as the empirical outcome. Appendix 1.D gives the aggregation formulas and numerical implementation.

1.6.5 Estimates and decomposition

This subsection reports the structural parameter estimates, model fit, and decomposition.

Table 1.6 reports the eight estimated behavioural parameters. The remaining inputs are calibrated directly from the pre-reform aggregate moments and the wage-experiment context.

Table 1.7 reports the aggregate moments most directly related to the wage experiment and compares the targeted DiD moments with their model counterparts. The model closely reproduces the pre-reform application rates and the wage-induced changes in applications and applicant quality.

Figure 1.10 presents the main decomposition as an accounting identity. The equilibrium effect on A and the contribution from spillovers onto the comparison group, $-\Delta Y_B$, sum to the structural DiD. The lower box reports the empirical DiD supplied to the estimator.

Table 1.8 reports the same decomposition and splits the B response into its two structural channels: the continuation-value effect and the hiring-probability effect.

Table 1.6: Structural parameter estimates

Parameter	Interpretation	Estimate
b	Baseline value of continuing search	0.1637
b_1	Slope of continuing-search value in average quality rank	0.6737
c_A	Application cost, occupation A	0.1023
c_B	Application cost, occupation B	0.0875
s_A	Worker payoff dispersion, occupation A	0.4976
s_B	Worker payoff dispersion, occupation B	0.3106
σ_A	Employer-choice dispersion, occupation A	0.7879
σ_B	Employer-choice dispersion, occupation B	0.7366

Notes: The table reports point estimates from the current GMM output. Sampling standard errors are not reported in this version.

Table 1.7: Model fit and wage-change moments

Moment	Data	Model	Standardized residual
Pre-reform applications per vacancy, A	10.270387	9.626326	1.496912
Pre-reform applications per vacancy, B	7.735478	7.481372	0.817640
DiD: applications per vacancy	0.210400	0.220075	-0.437778
DiD: average applicant quality	0.009500	0.003960	1.629384


Notes: Applications are in levels and applicant quality is an occupation-specific percentile rank. The overidentification statistic is $J = 710.325$ with 20 degrees of freedom and p -value 1.45×10^{-137} .

For applications, the model-based difference-in-differences effect of 0.214 log points comprises a 0.200-log-point increase in occupation A and a 0.013-log-point decline in occupation B. Comparison-group displacement therefore accounts for about 6% of the contrast, while the equilibrium response in the treated occupation accounts for the remaining 94%. Nearly all of the response in B operates through the hiring-probability channel. For average applicant quality, the response in B is close to zero, so the model-based difference-in-differences effect is almost entirely attributable to the improvement in occupation A.

1.6. STRUCTURAL DECOMPOSITION

45

Figure 1.10: Structural decomposition of the DiD



Notes: The control-group component equals minus the model response in occupation B , so a decline in B contributes positively to the DiD. The arrow links the structural DiD to its empirical counterpart.

Table 1.8: Decomposition of the reduced-form estimators

<i>Decomposition of the DiD effect</i>				
Outcome	Empirical DiD	Model DiD	Effect on A	Effect on B
Applications per job	0.210400	0.213645	0.200175	-0.013470
Average app quality	0.009500	0.003992	0.004050	0.000058
<i>Decomposition of the effect on B</i>				
Outcome	Continuation value		Hiring prob.	Adding-up err.
Applications per job	-0.000324		-0.013146	4.90×10^{-10}
Average app quality	-0.000012		0.000070	7.61×10^{-13}

Notes: The model DiD equals the effect on A minus the effect on B . Within B , the continuation-value and hiring-probability channels sum to the total B effect up to numerical path-integration error.

1.7 Conclusion

We study an occupation-wide increase in the default wage on a platform for short-term hospitality work. Raising the waiting-staff wage from €16 to €17 increased applications by approximately 23% relative to three neighbouring occupations, implying a wage elasticity of about 3.5. Mean applicant quality rose modestly, with an elasticity of approximately 0.3, and the improvement was concentrated at the bottom of the applicant-quality distribution. Hires increased by about 11%, implying an elasticity of 1.8. Together, these estimates show that occupation-wide wage changes can affect not only how many workers apply, but also the composition of the applicant pool and the probability that employers fill their vacancies.

The platform setting also makes clear why a conventional difference-in-differences estimate need not equal the equilibrium response in the treated occupation. Workers receive opportunities in several occupations, so a higher wage in one occupation can divert applications from another. Our structural decomposition shows that this spillover is modest. The decline in applications to the comparison occupations accounts for about 6% of the model-based application contrast, while the remaining 94% reflects the equilibrium response in the treated occupation. The spillover in applicant quality is economically negligible. The reduced-form elasticities therefore closely approximate the equilibrium responses of waiting-staff vacancies, and the conclusion that the higher wage attracted substantially more and somewhat better applicants is robust to interference across occupations.

References

- Altmann, Steffen, Anita Glenney, Robert Mahlstedt, and Alexander Sebald (2022). "The direct and indirect effects of online job search advice". In.
- Azar, José, Emiliano Huet-Vaughn, Ioana Marinescu, Bledi Taska, and Till von Wachter (2024). "Minimum Wage Employment Effects and Labor Market Concentration". In: *Review of Economic Studies* 91.4, pp. 1843–1883.
- Balgova, Maria, Tsegay Tekleselassie, Lukas Hensel, and Marc J. Witte (2025). *Wage Information and Applicant Selection*. IZA Discussion Paper 18220. Institute of Labor Economics (IZA).
- Banfi, Stefano and Benjamín Villena-Roldán (July 2019). "Do High-Wage Jobs Attract More Applicants? Directed Search Evidence from the Online Labor Market". In: *Journal of Labor Economics* 37.3, pp. 715–746. ISSN: 0734-306X. DOI: 10.1086/702627. URL: <https://www.journals.uchicago.edu/doi/10.1086/702627> (visited on 03/29/2024).
- Bassier, Ihsaan (2022). *Collective Bargaining and Spillovers in Local Labor Markets*. CEP Discussion Paper 1895. Centre for Economic Performance, London School of Economics and Political Science.
- Bassier, Ihsaan, Arindrajit Dube, and Suresh Naidu (2022). "Monopsony in Movers: The Elasticity of Labor Supply to Firm Wage Policies". In: *Journal of Human Resources* 57.S, S50–S86. DOI: 10.3368/jhr.monopsony.0319-10111R1.
- Bassier, Ihsaan, Alan Manning, and Barbara Petrongolo (2025). "Vacancy Duration and Wages". In: *Review of Economics and Statistics*, pp. 1–28. DOI: 10.1162/rest_a_01580.
- Beerli, Andreas, Stefano Fiorin, Andreas Gulyas, Daniel Kopp, Mahsa Khoshnama, and Michael Siegenthaler (2026). *The Role of Wages and Job Benefits in Job Search: Evidence from a Large-Scale Online Field Experiment*. URL: https://population-economics.committee.socialpolitik.de/sites/default/files/2025-01/Siegenthaler%20Value_of_Job_Benefits%20%285%29_0.pdf (visited on 06/16/2026).
- Belot, Michele, Philipp Kircher, and Paul Muller (2019). "Providing advice to job-seekers at low cost: An experimental study on online advice". In: *The review of economic studies* 86.4, pp. 1411–1447.
- Belot, Michèle, Philipp Kircher, and Paul Muller (2022). "How Wage Announcements Affect Job Search: A Field Experiment". In: *American Economic Journal: Macroeconomics* 14.4, pp. 1–67. DOI: 10.1257/mac.20200116.
- (2026). "Do the long-term unemployed benefit from automated occupational advice during online job search?" In: *The Economic Journal* 136.673, pp. 184–206.

Belot, Michèle, Bart K de Koning, Didier Fouarge, Philipp Kircher, Paul Muller, and Sandra Phippen (2025). *Advising job seekers in occupations with poor prospects: A field experiment*. Tech. rep. National Bureau of Economic Research.

Caldwell, Sydnee and Oren Danieli (2024). "Outside Options in the Labour Market". In: *Review of Economic Studies* 91.6, pp. 3286–3315. DOI: [10.1093/restud/rdae006](https://doi.org/10.1093/restud/rdae006).

Card, David (Apr. 2022). "Who Set Your Wage?" en. In: *American Economic Review* 112.4, pp. 1075–1090. ISSN: 0002-8282. DOI: [10.1257/aer.112.4.1075](https://doi.org/10.1257/aer.112.4.1075). URL: <https://www.aeaweb.org/articles?id=10.1257/aer.112.4.1075> (visited on 08/03/2023).

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer (2019). "The Effect of Minimum Wages on Low-Wage Jobs". In: *Quarterly Journal of Economics* 134.3, pp. 1405–1454. DOI: [10.1093/qje/qjz014](https://doi.org/10.1093/qje/qjz014).

Clemens, Jeffrey, Lisa B. Kahn, and Jonathan Meer (2021). "Dropouts Need Not Apply? The Minimum Wage and Skill Upgrading". In: *Journal of Labor Economics* 39.S1, S107–S149. DOI: [10.1086/711490](https://doi.org/10.1086/711490).

Dal Bó, Ernesto, Frederico Finan, and Martín A. Rossi (Aug. 2013). "Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service". In: *The Quarterly Journal of Economics* 128.3, pp. 1169–1218. ISSN: 0033-5533. DOI: [10.1093/qje/qjt008](https://doi.org/10.1093/qje/qjt008). URL: <https://doi.org/10.1093/qje/qjt008> (visited on 03/29/2024).

Datta, Nikhil (2023). *The Measure of Monopsony: The Labour Supply Elasticity to the Firm and Its Constituents*. CEP Discussion Paper 1930. Centre for Economic Performance, London School of Economics and Political Science.

Demir, Gizem Kosar (2024). *Labor Market Frictions and Spillover Effects from Publicly Announced Sectoral Minimum Wages*. IZA Discussion Paper 17510. Institute of Labor Economics (IZA).

Derenoncourt, Ellora and Claire Montialoux (2021). "Minimum Wages and Racial Inequality". In: *Quarterly Journal of Economics* 136.1, pp. 169–228. DOI: [10.1093/qje/qjaa031](https://doi.org/10.1093/qje/qjaa031).

Derenoncourt, Ellora and David Weil (2025). "Voluntary Minimum Wages: The Local Labor Market Effects of National Retailer Policies". In: *Quarterly Journal of Economics* 140.3, pp. 1901–1958. DOI: [10.1093/qje/qjaf016](https://doi.org/10.1093/qje/qjaf016).

Dube, Arindrajit, Jeff Jacobs, Suresh Naidu, and Siddharth Suri (2020). "Monopsony in Online Labor Markets". In: *American Economic Review: Insights* 2.1, pp. 33–46. DOI: [10.1257/aeri.20180150](https://doi.org/10.1257/aeri.20180150).

Dustmann, Christian, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp vom Berge (2022). "Reallocation Effects of the Minimum Wage". In: *Quarterly Journal of Economics* 137.1, pp. 267–328. DOI: [10.1093/qje/qjab028](https://doi.org/10.1093/qje/qjab028).

REFERENCES

49

- 62 Giupponi, Giulia, Robert Joyce, Attila Lindner, Tom Waters, Tom Wernham, and Xiaowei Xu (2024). "The Employment and Distributional Impacts of Nationwide Minimum Wage Changes". In: *Journal of Labor Economics* 42.S1, S293–S333. DOI: 10.1086/728471.
- 40 Giupponi, Giulia and Stephen J. Machin (2026). *Company Wage Policy in a Low-Wage Labor Market*. CEPR Discussion Paper 21345. Centre for Economic Policy Research.
- 27 Gopalan, Radhakrishnan, Barton H. Hamilton, Ankit Kalda, and David Sovich (2021). "State Minimum Wages, Employment, and Wage Spillovers: Evidence from Administrative Payroll Data". In: *Journal of Labor Economics* 39.3, pp. 673–707. DOI: 10.1086/711355.
- 29 Harasztosi, Péter and Attila Lindner (2019). "Who Pays for the Minimum Wage?" In: *American Economic Review* 109.8, pp. 2693–2727. DOI: 10.1257/aer.20171445.
- 3 Hirsch, Boris, Elke J. Jahn, Alan Manning, and Michael Oberfichtner (2022). *The Wage Elasticity of Recruitment*. IZA Discussion Paper 15675. Institute of Labor Economics (IZA).
- 99 Jäger, Simon, Christopher Roth, Nina Roussille, and Benjamin Schoefer (2024). "Worker beliefs about outside options". In: *The Quarterly Journal of Economics* 139.3, pp. 1505–1556.
- 35 Klaeui, Jeremias (2026). "When to Broaden, When to Focus: Job Search and Market Tightness". In.
- 6 Manning, Alan (Jan. 2021). "Monopsony in Labor Markets: A Review". en. In: *ILR Review* 74.1, pp. 3–26. ISSN: 0019-7939. DOI: 10.1177/0019793920922499. URL: <https://doi.org/10.1177/0019793920922499> (visited on 08/21/2023).
- 3 Marinescu, Ioana and Ronald Wolthoff (Apr. 2020). "Opening the Black Box of the Matching Function: The Power of Words". In: *Journal of Labor Economics* 38.2, pp. 535–568. ISSN: 0734-306X. DOI: 10.1086/705903. URL: <https://www.journals.uchicago.edu/doi/abs/10.1086/705903> (visited on 08/03/2023).
- 59 Matsudaira, Jordan D. (2014). "Monopsony in the Low-Wage Labor Market? Evidence from Minimum Nurse Staffing Regulations". In: *Review of Economics and Statistics* 96.1, pp. 92–102.
- 25 Mueller, Andreas I, Damian Osterwalder, Josef Zweimüller, and Andreas Kettmann (May 2023). "Vacancy Durations and Entry Wages: Evidence from Linked Vacancy–Employer–Employee Data". In: *The Review of Economic Studies*. ISSN: 0034-6527. DOI: 10.1093/restud/rdad051. URL: <https://doi.org/10.1093/restud/rdad051> (visited on 08/16/2023).
- 140 Ransom, Michael R. and David P. Sims (2010). "Estimating the Firm's Labor Supply Curve in a "New Monopsony" Framework: Schoolteachers in Missouri". In: *Journal of Labor Economics* 28.2, pp. 331–355. DOI: 10.1086/649904.

Sokolova, Anna and Todd Sorensen (Jan. 2021). "Monopsony in Labor Markets: A Meta-Analysis". en. In: *ILR Review* 74.1, pp. 27–55. ISSN: 0019-7939. DOI: 10.1177/0019793920965562. URL: <https://doi.org/10.1177/0019793920965562> (visited on 08/21/2023).

Staiger, Douglas O., Joanne Spetz, and Ciaran S. Phibbs (2010). "Is There Monopsony in the Labor Market? Evidence from a Natural Experiment". In: *Journal of Labor Economics* 28.2, pp. 211–236. DOI: 10.1086/652734.

1.A. SCREENSHOTS OF THE PLATFORM

51

Appendix 1.A Screenshots of the platform

1.A.1 Firm user experience

Quel est le type de votre établissement ?

Cafés / Restaurants	Bars / Clubs	Hôtel
Vente à emporter	Retail / Vente	Santé
Restauration Collective	Événementiel	Autres

(a) Category of establishment choice

Décrivez brièvement votre établissement

Lieu chaleureux, clientèle de quartier. Existe à Pigalle depuis 1899

0 / 2000

Les établissements ayant une description attrayante attirent 2 fois plus de candidats !

(b) Establishment description

Pour un poste en ...

Hôtel Cuisine Salle / Bar

(c) Qualification family choice

Pour un poste en ...

Hôtel Cuisine Salle / Bar

... En tant que

Chef de partie

Demi-chef de partie

Commis de cuisine

Second de cuisine

Chef de cuisine

Plongeur

→ Voir plus de catégories

(d) Qualification choice

Vous pourrez choisir votre extra favori directement lors de la création de votre Mission / Extra.

Mission / Extra CDI Saisonnier

(e) Type of contract choice

De combien d'extras avez-vous besoin ?

1

(f) Number of required workers

Octobre 2024

1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

Novembre 2024

1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30

Choisir une annonce

(g) Dates of job

ven. 12/11 Début 11:15 Fin 19:30 Pas de pause

sem. 16/11 Début 08:30 Fin 17:00 Pas de pause

dim. 17/11 Début 08:30 Fin 17:00 Pas de pause

lun. 18/11 Début 08:30 Fin 17:00 Pas de pause

mar. 19/11 Début 08:30 Fin 17:00 Pas de pause

mer. 20/11 Début 08:30 Fin 17:00 Pas de pause

jeu. 21/11 Début 08:30 Fin 17:00 Pas de pause

ven. 22/11 Début 08:30 Fin 17:00 Pas de pause

sam. 23/11 Début 08:30 Fin 17:00 Pas de pause

Valider ces horaires (84h45)

(h) Detailed schedule of job

Précisez les missions, les avantages et le profil recherché

Quelles sont les missions principales ? Y a-t-il des pourboires ? Quel est le niveau d'expérience attendu ? Quelle est l'ambiance de votre établissement ? Quel style de clientèle avez-vous ?

0 / 5000

Une annonce avec une description attire 2 fois plus de candidats !

(i) Job description

Soyez plus attractif

En prévision des JO, nous vous conseillons de booster vos demandes de +20% afin de réduire les refus et d'améliorer la qualité des profils reçus !

+10% +20% +30% +40% +50%

Personnaliser 0 %

Prestation

Durée totale 75000

Prix horaire HT 17,90 €

Total HT* 1 253,00 €

*Ce total inclut le prix horaire ?

Vous ne pouvez qu'une fois la prestation de candidat effectuée

Suivre

(j) Wage choice

Total HT 1 375,30 € (+20%)

Poste Plongeur

Nombre d'extras 1

Période du 12/11/2024 au 23/11/2024 Total 75000

Adresse de la mission rue des boulevards, Paris 75006

Tenue T-shirt Blanc Pantalon Jean Bleu Aligne tailleur Noir

Description Plongeur

Pублиer cette annonce

(k) Job ad summary

Actuellement en ligne

Offre boostée

Plongeur

5 postulants

1 extra

Du mar. 12/11 à 08:30 au 23/11 rue des boulevards 75006 Paris

Dupliquer Voir la demande

Afficher détails 7/11/2024

(l) Number of applicants to the ad

Laura

Serveur Maître d'hôtel restaurant

Vendeur en boulangerie / Cuisinier de cuisine / Receptionniste

Vendeur Boulangerie

91290, Saint-Sauveur

29 ans

Expériences

Serveur 1 an Paris, restaurant Extra

Maître d'hôtel restaurant 1 an Paris, Boulangerie Côté Extra, mai 2022

Sélectionner ce candidat

(m) Applicant profile

Lisa

Barmen

J'ai 19 ans, j'ai travaillé comme barmen dans la salle brasserie la semaine à côté de la station métro pendant 3 mois et j'aimerais dans le 12 ème mais je n'ai pas beaucoup d'expérience avec le plateau. Mais je suis motivé.

75012, Paris

27 ans

Expériences

Barmen La brasserie, salle brasserie Extra

Bien, elle est dynamique et vient 1 jour par semaine nous aider, on a un bar d'encre surtout pour les happy hours.

Connaissances

Anglais

Sélectionner ce candidat

(n) Applicant profile

Félicitations !

Le candidat va recevoir une notification

La suite ?


Nous vous conseillons de prendre contact avec l'extra pour vous assurer de sa venue

Contactez l'extra


(o) Applicant selection

CHAPTER 1. THE WAGE ELASTICITY OF JOB APPLICATIONS


1.A.2 Worker user experience




(a) Job ad example




(b) Job ad summary



(c) Establishment description



(d) Application final step



(e) Application tracking

Appendix 1.B Estimating quality

We use employers' revealed choices to construct a quality score for each worker–job pair. Let H_{ij} indicate that worker i is hired for job j , conditional on applying. We estimate a class-weighted logistic model and define the fitted score as

$$S_{ij} = \frac{1}{1 + \exp(-z_{ij})}, \quad z_{ij} = \beta_0 + X'_{ij}\beta + \gamma_{c(j),e(j)}. \quad (1.33)$$

The vector X_{ij} contains worker-, job-, and pair-level characteristics, and $\gamma_{c(j),e(j)}$ denotes establishment fixed effects indexed by establishment category and employer. Worker characteristics include gender, age quintile, average rating, and time since the last connection to the platform. The job controls include the number of workers requested. Pair-level characteristics include quintiles of straight-line distance between the worker's residence and the workplace; indicators for living in the same postcode or *département*; whether the worker has previously worked for and been approved by the employer; and detailed measures of previous experience. We group experience into the most recent assignment, the second through fifth most recent assignments, and the sixth through tenth most recent assignments. For each group, we include average ratings and indicators for whether the experience was obtained through the platform, was in the same occupation or establishment category as the current job, or was in a higher, lower, or undefined occupation within the same family. We rank occupations within families by their mean posted hourly wage. Missing-value indicators are included for age, average ratings, experience ratings, and time since the last connection.

Hires are rare, so the estimation weights hired and non-hired observations inversely by their respective sample frequencies. This weighting improves the model's ability to rank hires but changes the calibration of its fitted values. We therefore interpret S_{ij} as a quality score, not as the applicant's actual hiring probability.

We randomly assign 20% of service-sector jobs posted between 1 January 2023 and 30 December 2025 to the training sample and use the remaining 80% as the test sample. The broader service-sector data contain 22 occupations. Eighteen are not suitable comparison occupations for the event study, but they provide useful information about employers' hiring choices. Before estimation-specific exclusions, the training sample contains 382,201 worker–job pairs, 10,989 workers, and 23,165 jobs; the test sample contains 4,155,162 pairs, 20,708 workers, and 237,945 jobs. The effective logistic-regression sample contains 361,689 observations after fixed-effect and missing-value handling. Quality is pair-specific: a worker can receive different scores for different jobs.

Table A1 compares worker–job pairs in the training, test, and DiD samples. Although the test group contains 80% of jobs, it contains 92% of pairs because applicant-pool sizes vary substantially across postings. Random assignment makes the training and test samples broadly comparable in worker demographics, geographic proximity, prior worker–employer relationships, ratings, job requirements, and work experience. The DiD sample is instead determined by the event-study period and four selected occupations. It consequently contains more women, younger workers, and somewhat different occupational histories. We obtain the quality outcome used in the DiD analysis by applying the model estimated on the training sample to this analysis sample.

1.B. ESTIMATING QUALITY

55

Table A1: Summary statistics of worker-job pairs across samples

	Training Sample	Test Sample	DiD Sample
Female	0.44 (0.50)	0.37 (0.48)	0.49 (0.50)
Age	29.37 (7.15)	30.21 (7.69)	28.15 (6.47)
Same postcode	0.02 (0.13)	0.02 (0.12)	0.02 (0.13)
Same department	0.32 (0.47)	0.29 (0.45)	0.31 (0.46)
Prior worker-establishment history	0.08 (0.27)	0.06 (0.24)	0.07 (0.25)
Avg. rating	4.59 (0.52)	4.58 (0.53)	4.53 (0.58)
Nb. of extras demanded	1.32 (1.20)	1.21 (1.13)	1.33 (1.38)
Crow distance (km)	15.68 (41.37)	16.55 (43.12)	17.27 (48.26)
Experience (most recent)			
Avg. rating	4.79 (0.83)	4.81 (0.80)	4.80 (0.82)
Same qualification	0.33 (0.47)	0.37 (0.48)	0.35 (0.48)
Same category	0.12 (0.33)	0.15 (0.35)	0.10 (0.30)
Rank	13.79 (4.79)	14.59 (5.66)	14.71 (4.81)
Same Family, higher rank	0.27 (0.44)	0.21 (0.41)	0.13 (0.33)
Same Family, lower rank	0.20 (0.40)	0.16 (0.37)	0.25 (0.44)
Different family	0.20 (0.40)	0.26 (0.44)	0.27 (0.44)
Experience (last 2-5)			
Avg. rating	4.65 (0.94)	4.66 (0.94)	4.62 (0.97)
Same qualification	0.34 (0.48)	0.37 (0.48)	0.37 (0.48)
Same category	0.12 (0.32)	0.14 (0.35)	0.10 (0.30)
Rank	13.84 (4.69)	14.64 (5.64)	14.73 (4.78)
Same Family higher rank	0.28 (0.45)	0.21 (0.41)	0.12 (0.32)
Same Family lower rank	0.18 (0.39)	0.16 (0.37)	0.24 (0.43)
Different family	0.18 (0.39)	0.25 (0.43)	0.25 (0.43)
Experience (last 6-10)			
Avg. rating	4.69 (0.84)	4.70 (0.85)	4.66 (0.90)
Same qualification	0.30 (0.46)	0.33 (0.47)	0.31 (0.46)
Same category	0.10 (0.30)	0.12 (0.33)	0.08 (0.28)
Rank	12.68 (5.91)	13.10 (6.94)	13.14 (6.46)
Same Family higher rank	0.26 (0.44)	0.20 (0.40)	0.11 (0.31)
Same Family lower rank	0.17 (0.37)	0.14 (0.35)	0.23 (0.42)
Different family	0.18 (0.38)	0.22 (0.41)	0.23 (0.42)
N (worker-job pairs)	382,201	4,155,162	320,833
Share of sample	0.08	0.92	0.07

Notes: The table reports summary statistics for worker-job pairs in the training, test, and DiD samples. Standard deviations are in parentheses. The DiD sample contains applications to waiting-staff, bar-staff, receptionist, and runner jobs posted between 5 April 2023 and 4 May 2024.

Table A2 reports the estimated coefficients for the variables in the service-sector logistic model. These raw logit coefficients are descriptive of the fitted model; because the regressors use different scales, their magnitudes should not be interpreted as a formal variable-importance ranking.

Table A3 evaluates the quality model in the test sample. An AUC of 0.76 means that the model assigns a higher score to a randomly selected hired applicant than to a randomly selected non-hired applicant 76% of the time. At the threshold of 0.6, which maximises Youden’s J , sensitivity is 0.605 and specificity is 0.785. Thus Youden’s $J = 0.605 + 0.785 - 1 = 0.390$. Precision is 0.15, reflecting the rarity of hiring: only approximately 3.7% of test-sample applications result in a hire. These metrics show that the score contains meaningful ranking information, while also illustrating the difficulty of classifying rare outcomes.


Table A3: Prediction evaluation metrics

AUC	Sensitivity	Specificity	Accuracy	Precision	Youden Index	Threshold
0.76	0.605	0.785	0.774	0.15	1.39	0.6

Notes: The table reports classification metrics for the applicant-quality model in the test sample. The threshold maximises Youden’s J , defined as sensitivity plus specificity minus one.

Figure A3 plots the receiver operating characteristic (ROC) curve. The curve lies above the 45-degree line associated with random classification. Its shape illustrates the trade-off between sensitivity and specificity: lowering the classification threshold increases both the true-positive and false-positive rates. The AUC summarises the model’s ability to rank applicants across all possible thresholds.

Figure A3: Receiver operating characteristic curve



Notes: The curve plots sensitivity against one minus specificity across all classification thresholds in the test sample. Table A3 reports the corresponding evaluation metrics.

1.B. ESTIMATING QUALITY

57

Table A2: Coefficients from the quality model

Variable	Estimate	Std. Error	Variable	Estimate	Std. Error
Male dummy	-0.207	0.006	Last 6-10 experiences rank	-0.005	0.001
Same postcode	0.212	0.022	Best experience internal	0.003	0.007
Same department	0.124	0.008	Best experience from the same family, better paid occupation	0.007	0.040
Nb. of extras demanded	0.236	0.003	Best experience from the same family, less well paid occupation	0.227	0.040
Avg. rating	0.303	0.007	Best experience rating missing	-0.120	0.060
Avg. rating missing	0.755	0.135	Best experience from the same category	0.042	0.011
Prior work history	2.413	0.009	Best experience from the same occupation	0.238	0.040
Most recent experience missing	1.434	0.077	Best experience from a different family	-0.150	0.040
Most recent experience internal	0.198	0.015	Reference category		
Most recent experience rating	0.044	0.004	Age (quintile 1)	0.064	0.008
Most recent experience rating missing	-0.014	0.006	Age (quintile 2)	-0.004	0.009
Most recent experience from the same occupation	0.271	0.011	Age (quintile 3)	-0.017	0.009
Most recent experience from the same category	0.358	0.009	Age (quintile 4)	-0.144	0.010
Most recent experience from the same family, better paid occupation	0.068	0.013	Age (quintile 5)	-0.175	0.100
Most recent experience from the same family, less well paid occupation	0.271	0.012	Reference category		
Last 2-5 experiences missing	-0.288	0.051	Crow distance (quintile 1)	-0.069	0.009
Last 2-5 experiences internal	0.252	0.021	Crow distance (quintile 2)	-0.069	0.011
Last 2-5 experiences rating	0.092	0.004	Crow distance (quintile 3)	-0.111	0.011
Last 2-5 experiences rating missing	-0.020	0.008	Crow distance (quintile 4)	-0.208	0.012
Last 2-5 experiences from the same qualification	0.276	0.012	Crow distance (quintile 5)	Reference category	
Last 2-5 experiences from the same category	0.192	0.010	Connected over a week ago	0.062	0.038
Last 2-5 experiences from the same family, better paid occupation	0.106	0.014	Connected in last week	0.225	0.037
Last 2-5 experiences from the same family, less well paid occupation	0.136	0.013	Connected in the last day	0.333	0.037
Last 6-10 experiences missing	0.008	0.026	Connected in the 12 hours	0.342	0.037
Last 6-10 experiences internal	0.015	0.014	Connected in the 6 hours	0.510	0.036
Last 6-10 experiences rating	0.114	0.004	Connected in the last hour	-9.248	19.634
Last 6-10 experiences rating missing	-0.023	0.008	Last connection missing	Reference category	
Last 6-10 experiences from the same qualification	0.258	0.011	English level: none	0.288	0.017
Last 6-10 experiences from the same category	0.099	0.010	English level: missing	0.050	0.019
Last 6-10 experiences from the same family, better paid occupation	-0.008	0.013	English level: beginner	0.222	0.015
Last 6-10 experiences from the same family, less well paid occupation	0.225	0.012	English level: intermediate	0.304	0.016
Last experience rank	-0.003	0.001	English level: proficient	0.306	0.015
Last 2-5 experiences rank	-0.004	0.001	English level: advanced		

Notes: The table reports the coefficients and standard errors estimated from equation (1.33).

Appendix 1.C Conceptual framework: derivations and equilibrium

This appendix derives the worker's application rule, defines the exact Poisson-logit equilibrium, relates it to the mean-field representation used in the main text, and proves Proposition 1.1. Throughout, F is the stationary joint distribution of types among workers who are unemployed and available on the platform, and F_j^N is the notification-conditional distribution defined in equation (1.2).

1.C.1 Worker problem

The timing stated in the main text is important for the discount factors below. An unemployed worker receives b during the current period. A notification and the application decision then occur, the application cost is paid immediately, and a successful applicant starts employment next period. The match-specific shock remains constant during the employment spell. Let $W_j(\varepsilon_j; q, w)$ be the value at the beginning of an employment period in occupation j . Employment ends at the close of a period with probability δ , so

$$W_j(\varepsilon_j; q, w) = w_j + \varepsilon_j + \tilde{\beta} [(1 - \delta)W_j(\varepsilon_j; q, w) + \delta U(q; w)]. \quad (1.34)$$

It follows that

$$W_j(\varepsilon_j; q, w) - U(q; w) = \frac{w_j + \varepsilon_j - (1 - \tilde{\beta})U(q; w)}{1 - \tilde{\beta}(1 - \delta)}. \quad (1.35)$$

At the beginning of an unemployment period, the worker's value satisfies

$$U(q; w) = b + \tilde{\beta} U(q; w) + \sum_{j \in \{A, B\}} \alpha_j(q) \mathbb{E}_{\varepsilon_j} \left[\max \left\{ 0, -\tilde{c}_j + \tilde{\beta} p_j(q; w) [W_j(\varepsilon_j; q, w) - U(q; w)] \right\} \right], \quad (1.36)$$

The term $\tilde{\beta} p_j(W_j - U)$ reflects that a successful applicant enters employment next period, whereas an unsuccessful applicant begins that period unemployed. Define

$$u = (1 - \tilde{\beta})U, \quad \beta = \frac{\tilde{\beta}}{1 - \tilde{\beta}(1 - \delta)}, \quad c_j = \frac{\tilde{c}_j}{\beta},$$

where β is the discounted employment-spell multiplier for employment that be-

1.C. CONCEPTUAL FRAMEWORK: DERIVATIONS AND EQUILIBRIUM 59

gins next period. Substitution of (1.35) into (1.36) gives

$$-\tilde{c}_j + \tilde{\beta} p_j [W_j - U] = \beta \{ p_j(q_j; w) [w_j + \varepsilon_j - u(q; w)] - c_j \}. \quad (1.37)$$

Because $\beta > 0$, the application rule is

$$p_j(q_j; w) [w_j + \varepsilon_j - u(q; w)] \geq c_j.$$

1.C.2 Application threshold

Rearranging the application rule yields the threshold

$$\varepsilon_j \geq \bar{\varepsilon}_j(q; w) \equiv \frac{c_j}{p_j(q_j; w)} - [w_j - u(q; w)].$$

For a logistic shock with scale s_j ,

$$\Pr(\varepsilon_j \geq x) = \frac{1}{1 + \exp(x/s_j)}.$$

Evaluating this survival function at $\bar{\varepsilon}_j(q; w)$ proves equation (1.4). Its index is increasing in w_j and p_j , holding u fixed, which establishes the stated monotonicity.

The value of the option to apply can also be written in terms of the application probability. For any integrable random variable, the tail-integral formula gives

$$\mathbb{E}[(\varepsilon_j - \bar{\varepsilon}_j)^+] = \int_{\bar{\varepsilon}_j}^{\infty} \Pr(\varepsilon_j > x) dx. \quad (1.38)$$

For the logistic survival function,

$$\begin{aligned} \int_{\bar{\varepsilon}_j}^{\infty} \frac{1}{1 + \exp(x/s_j)} dx &= s_j \log \left[1 + \exp \left(-\frac{\bar{\varepsilon}_j}{s_j} \right) \right] \\ &= -s_j \log(1 - \lambda_j), \end{aligned} \quad (1.39)$$

where the last equality uses

$$1 - \lambda_j = \frac{1}{1 + \exp(-\bar{\varepsilon}_j/s_j)}.$$

Equation (1.37) also implies that the realised value of receiving a j -notification is

$$\beta p_j(q_j; w) (\varepsilon_j - \bar{\varepsilon}_j)^+.$$

Rearranging the unemployment Bellman equation and applying (1.39) therefore gives

$$u(q; w) = b - \sum_{j \in \{A, B\}} \alpha_j(q) \beta s_j p_j(q; w) \log[1 - \lambda_j(q; w)],$$

which is equation (1.5).

1.C.3 Worker policy functions

Fix wages w and hiring probabilities p_A, p_B . For each worker type, the continuation value is the solution to

$$\Psi(u; p, w) \equiv u - b + \sum_{j \in \{A, B\}} \alpha_j(q) \beta s_j p_j \log \left[1 - \frac{1}{1 + \exp\{-[w_j - u - c_j/p_j]/s_j\}} \right] = 0. \quad (1.40)$$

Because

$$\frac{\partial \Psi}{\partial u} = 1 + \sum_{j \in \{A, B\}} \alpha_j(q) \beta p_j \lambda_j > 0,$$

the worker problem has at most one solution. Its limits as $u \rightarrow -\infty$ and $u \rightarrow +\infty$ have opposite signs, so a solution exists. Denote it by $m(q; p_A, p_B, w)$. The worker policy functions are

$$\ell_j(q; p_A, p_B, w) = \frac{1}{1 + \exp\{-[w_j - m(q; p_A, p_B, w) - c_j/p_j]/s_j\}}.$$

1.C.4 Exact Poisson-logit hiring operator

Let $\mathcal{Q} = [0, 1]^2$ and

$$a_j(q_j) = \exp(q_j/\sigma_j)$$

be the employer-choice index. Notifications sent by a vacancy form a Poisson process on $\mathcal{Q}$ with mean M_j and mark distribution F_j^N . Independent application decisions thin this process, so applications form a marked Poisson process with intensity measure

$$\mu_j(dq; w) = M_j \lambda_j(q; w) dF_j^N(q) = \frac{M_j}{\bar{\alpha}_j} \alpha_j(q) \lambda_j(q; w) dF(q). \quad (1.41)$$

Conditional on a focal worker applying with productivity q_j , the logit probability of being selected from a realised applicant pool is

$$\frac{a_j(q_j)}{a_j(q_j) + \sum_{i \in A_j} a_j(q_{ij})}.$$

1.C. CONCEPTUAL FRAMEWORK: DERIVATIONS AND EQUILIBRIUM 61

Under the Palm distribution of a Poisson process, conditioning on the focal application leaves the process of other applications unchanged. Hence, if

$$S_j(w) = \sum_{i \in \mathcal{A}_j} a_j(q_{ij}),$$

then the Laplace functional of the marked Poisson process is

$$\mathbb{E}[\exp\{-tS_j(w)\}] = \exp\left\{-M_j \int_{\mathcal{Q}} [1 - \exp\{-ta_j(x_j)\}] \lambda_j(x; w) dF_j^N(x)\right\}. \quad (1.42)$$

Using

$$\frac{a}{a+S} = \int_0^\infty a \exp[-t(a+S)] dt,$$

the probability of being hired conditional on applying is

$$p_j^E(q_j; w) = \kappa_j \int_0^\infty a_j(q_j) \exp[-ta_j(q_j)] \times \exp\left\{-M_j \int_{\mathcal{Q}} [1 - \exp\{-ta_j(x_j)\}] \lambda_j(x; w) dF_j^N(x)\right\} dt. \quad (1.43)$$

An exact equilibrium is a collection $(p_A^E, p_B^E, \lambda_A^E, \lambda_B^E, u^E)$ satisfying

$$\lambda_j^E(q; w) = \ell_j(q; p_A^E(q_A; w), p_B^E(q_B; w), w), \quad (1.44)$$

$$p_j^E(q_j; w) = \mathcal{P}_j^E[\lambda_j^E(\cdot; w)](q_j), \quad (1.45)$$

where $\mathcal{P}_j^E$ is the operator in (1.43); u^E is the unique solution of the worker problem conditional on p_A^E, p_B^E . The conceptual framework assumes that the resulting equilibrium is unique. More generally, the analysis requires a stable equilibrium-selection rule.

1.C.5 Mean-field approximation

The exact operator requires a one-dimensional Laplace integral for each occupation and productivity point. The mean-field approximation replaces the random competitor index $S_j(w)$ by

$$D_j(w) = \mathbb{E}[S_j(w)] = M_j \int_{\mathcal{Q}} a_j(x_j) \lambda_j(x; w) dF_j^N(x),$$

$$p_j^{MF}(q_j; w) = \kappa_j \frac{a_j(q_j)}{a_j(q_j) + D_j(w)}. \quad (1.46)$$

This is the representation in equations (1.7) and (1.8). It preserves smooth employer choice, productivity sorting, and the endogenous response of hiring probabilities

to the applicant pool while substantially reducing computation.

The Jensen inequality in the main text follows directly. For a fixed focal index $a > 0$, let $f_a(s) = a/(a + s)$. Since

$$f_a''(s) = \frac{2a}{(a + s)^3} > 0,$$

$$p_j^E(q_j; w) = \kappa_j \mathbb{E}[f_{a_j(q_j)}(S_j)] \geq \kappa_j f_{a_j(q_j)}(\mathbb{E}[S_j]) = p_j^{MF}(q_j; w).$$

The direction of the approximation is therefore unambiguous.

A second-order expansion around $D_j = \mathbb{E}[S_j]$ quantifies the Jensen gap:

$$p_j^E(q_j; w) = p_j^{MF}(q_j; w) + \kappa_j \frac{a_j(q_j) \text{Var}[S_j(w)]}{[a_j(q_j) + D_j(w)]^3} + R_{3j}(q_j; w), \quad (1.47)$$

$$\text{Var}[S_j(w)] = M_j \int_{\mathcal{Q}} a_j(x_j)^2 \lambda_j(x; w) dF_j^N(x). \quad (1.48)$$

The variance expression is the variance of a compound Poisson sum. The remainder R_{3j} collects terms involving third and higher centred moments; in particular $f_a'''(s) = -6a/(a + s)^4$. The approximation is accurate when the random competitor index is concentrated relative to $a_j(q_j) + D_j(w)$, not merely when its mean is large.

The limiting cases also differ across the exact and mean-field operators. As $\sigma_j \downarrow 0$, the realised employer choice approaches selection of the highest-productivity applicant. As $\sigma_j \rightarrow \infty$, all choice indices converge to one. Writing

$$\mathcal{A}_j(w) = M_j \int_{\mathcal{Q}} \lambda_j(q; w) dF_j^N(q),$$

the exact and mean-field limits are, respectively,

$$p_j^E(q_j; w) \rightarrow \kappa_j \frac{1 - \exp[-\mathcal{A}_j(w)]}{\mathcal{A}_j(w)}, \quad p_j^{MF}(q_j; w) \rightarrow \frac{\kappa_j}{1 + \mathcal{A}_j(w)}. \quad (1.49)$$

The first expression uses $\mathbb{E}[1/(1 + N)] = (1 - e^{-\mathcal{A}_j})/\mathcal{A}_j$ for $N \sim \text{Poisson}(\mathcal{A}_j)$; the second replaces N by its mean.

Finally, the exact operator satisfies the one-position adding-up identity

$$\int_{\mathcal{Q}} p_j^E(q_j; w) \mu_j(dq; w) = \kappa_j [1 - \exp\{-\mathcal{A}_j(w)\}], \quad (1.50)$$

because a nonempty applicant pool produces one selected applicant, who is successfully matched with probability κ_j . The pointwise mean-field lower bound generally does not preserve (1.50); quantities that depend only on whether the Poisson

applicant pool is empty are therefore computed from the exact Poisson model.

1.C.6 Proof of Proposition 1.1

Write equation (1.4) as

$$\lambda_j = L\left(\frac{w_j - u - c_j/p_j}{s_j}\right), \quad L'(z) = L(z)[1 - L(z)].$$

Taking the total derivative along an equilibrium wage path gives

$$\frac{d\lambda_j}{dw_k} = \frac{\lambda_j(1 - \lambda_j)}{s_j} \left[\mathbf{1}\{j = k\} - \frac{du}{dw_k} + \frac{c_j}{p_j^2} \frac{dp_j}{dw_k} \right],$$

which proves equation (1.10). The first term is positive for an own-wage change, while the other two terms need not be. Therefore the total derivative is not sign restricted. $\square$

For completeness, define

$$H_j(q) = -s_j \log[1 - \lambda_j(q)], \quad K_j(q) = H_j(q) + \frac{c_j}{p_j(q)} \lambda_j(q).$$

Differentiating the continuation-value equation and using the envelope property gives

$$\frac{du(q)}{dw_k} = \frac{\beta \alpha_k(q) p_k(q_k) \lambda_k(q) + \beta \sum_{j \in \{A, B\}} \alpha_j(q) K_j(q) \frac{dp_j(q_j)}{dw_k}}{1 + \beta \sum_{j \in \{A, B\}} \alpha_j(q) p_j(q_j) \lambda_j(q)}. \quad (1.51)$$

Holding hiring probabilities fixed, an increase in w_k strictly raises the continuation value for workers with a positive probability of applying to and obtaining a job in k . Hiring-probability adjustments can reinforce or attenuate this response.

Under the mean-field rule,

$$\frac{dp_j(q_j)}{dw_k} = -\frac{p_j(q_j)}{a_j(q_j) + D_j} \frac{dD_j}{dw_k}, \quad (1.52)$$

$$\frac{dD_j}{dw_k} = M_j \int_{\mathcal{Q}} a_j(x_j) \frac{d\lambda_j(x)}{dw_k} dF_j^N(x). \quad (1.53)$$

Equations (1.10), (1.51), and (1.52)–(1.53) form a linear fixed-point system for local equilibrium responses.

1.C.7 Applicant-pool outcomes

Applications to an occupation- j vacancy form the marked Poisson process with joint type-intensity measure $\mu_j(dq; w)$ defined in equation (1.41). The expected number of applications is $\mathcal{A}_j = \int_{\mathcal{Q}} \mu_j(dq)$, and the probability of a nonempty applicant pool is $1 - \exp(-\mathcal{A}_j)$. Conditional on a nonempty pool, mean quality is

$$\bar{Q}_j = \frac{\int_{\mathcal{Q}} q_j \mu_j(dq)}{\mathcal{A}_j}.$$

For minimum and maximum quality, let μ_j^Q be the push-forward of μ_j under the map $q \mapsto q_j$, and define

$$C_j(x) = \int_{[0,x]} \mu_j^Q(dv), \quad T_j(x) = \int_{[x,1]} \mu_j^Q(dv).$$

Conditional on a nonempty pool, expected minimum and maximum quality are

$$Q_{j,\min} = \int_0^1 \frac{\exp[-C_j(x)] - \exp[-\mathcal{A}_j]}{1 - \exp[-\mathcal{A}_j]} dx, \quad (1.54)$$

$$Q_{j,\max} = \int_0^1 \frac{1 - \exp[-T_j(x)]}{1 - \exp[-\mathcal{A}_j]} dx. \quad (1.55)$$

These expressions must be adapted if the empirical analysis assigns values to vacancies with no applications rather than conditioning on nonempty applicant pools.

Appendix 1.D Structural estimation and decomposition: additional details

This appendix describes the construction of the empirical moments, provides a formal identification argument for the mean-field model, and gives the GMM moment conditions and inference procedure. The baseline estimation uses empirical joint cells, cell-specific notification probabilities, and the mean-field hiring operator. An exact Poisson-logit re-estimation is reported as a sensitivity exercise. Gaussian-copula notation appears only as an optional sensitivity case.

1.D. STRUCTURAL ESTIMATION AND DECOMPOSITION

65

1.D.1 Empirical moment construction

Reference population and productivity cells

All baseline moments are constructed from the stable pre-reform period and refer to the same population of unemployed workers available on the platform and the same universe of vacancies and notifications. Worker quality is measured before the application decision. For the behavioral equations, we convert it into occupation-specific percentile ranks $q_A, q_B \in [0, 1]$, defined among unemployed available workers rather than in the selected applicant pool. Sample splitting or cross-fitting prevents a worker's realised application or hiring outcome for the same job from mechanically determining that worker's rank.

The reduced-form quality outcome remains the worker-job hiring score S_{ij} in its original units. Because it is worker-job specific, for each occupation and worker cell we average predicted scores over the relevant distribution of potential jobs, using information determined before application. Denote this cell-level score by $\chi_j(c)$. Thus $q_j(c)$ is the percentile-rank input to preferences and employer sorting, whereas $\chi_j(c)$ maps the same cell into the empirical quality-score outcome.

We partition each rank and cross the partitions. A joint cell c contains workers with similar (q_A, q_B) . Its weight ω_c is the share of unemployed available worker-periods in that cell before notification and application decisions. Using worker-periods is important because the notification probabilities are defined per model period. Applicant or hire shares are not valid cell weights because they are endogenous outcomes.

Conditional application and hiring moments

For each joint cell, the application moments are

$$\hat{\lambda}_A(c) = \frac{\#\{\text{A applications following A notifications in } c\}}{\#\{\text{A notifications in } c\}}, \quad (1.56)$$

$$\hat{\lambda}_B(c) = \frac{\#\{\text{B applications following B notifications in } c\}}{\#\{\text{B notifications in } c\}}. \quad (1.57)$$

The denominator is the relevant number of notifications, not the number of workers. Notifications with no application remain in the denominator.

The model implies $p_A = p_A(q_A)$ and $p_B = p_B(q_B)$. We therefore estimate

$$\hat{p}_A(x) = \frac{\#\{\text{hires from A applications in } q_A \text{ bin } x\}}{\#\{\text{A applications in } q_A \text{ bin } x\}}, \quad (1.58)$$

$$\hat{p}_B(x) = \frac{\#\{\text{hires from B applications in } q_B \text{ bin } x\}}{\#\{\text{B applications in } q_B \text{ bin } x\}}. \quad (1.59)$$

Thus hiring probabilities are conditional on applying. Hires divided by notifications combine p_j and λ_j and are not the relevant moment. We merge the one-dimensional hiring estimates back onto the joint cells for use in the application equations, but count each distinct rank-bin hiring estimate only once in the GMM criterion.

Aggregate moments

For occupation j , the aggregate pre-reform moments are

$$\hat{M}_j = \frac{\text{notifications sent by } j}{\text{vacancies in } j}, \quad \hat{A}_j = \frac{\text{applications received by } j}{\text{vacancies in } j}, \quad (1.60)$$

$$\hat{z}_j = \frac{\text{vacancies ending without a hire}}{\text{vacancies in } j}. \quad (1.61)$$

The corresponding worker-side notification probability is

$$\hat{\alpha}_j = \frac{\#\{j\text{-notifications}\}}{\#\{\text{unemployed available worker-periods}\}}. \quad (1.62)$$

For each joint cell c , we also compute

$$\hat{\alpha}_j(c) = \frac{j\text{-notifications sent to worker-periods in cell } c}{\text{unemployed available worker-periods in cell } c}.$$

The aggregate $\hat{\alpha}_j$ is a weighted average of these cell probabilities, but the structural equilibrium uses $\hat{\alpha}_j(c)$. The counts must satisfy

$$\hat{V}_j \hat{M}_j = \hat{N}^U \sum_c \hat{\omega}_c \hat{\alpha}_j(c), \quad (1.63)$$

where $\hat{V}_j$ is the number of occupation- j vacancies and $\hat{N}^U$ is the number of unemployed available worker-periods in the same observation window. We construct all objects from common underlying counts and report the residual in (1.63) as a data-reconciliation diagnostic.

Applications per vacancy are computed in a held-out aggregation fold as applications divided by vacancies. They are not derived from M_j , notification shares, and cell application rates estimated in the cell-moment fold. The equality

$$A_j = M_j \sum_c \pi_j(c) \lambda_j(c)$$

is consequently an overidentifying restriction rather than an identity imposed while constructing the empirical moments. If no held-out or otherwise independent aggregate moment is available, applications per vacancy are kept as a fit

1.D. STRUCTURAL ESTIMATION AND DECOMPOSITION

67

diagnostic and the mechanically redundant direction is excluded from the GMM criterion and the degrees of freedom of J . The time unit used to calculate notification probabilities is also used for match duration, discounting, and all flow payoffs. Standard errors account for repeated workers and common vacancy or time shocks using the clustering structure of the empirical analysis.

Representative comparison occupation

The empirical comparison group contains several non-A occupations. In the two-occupation approximation, we pool their notifications, applications, vacancies, and hires before constructing the B moments. We define the representative pre-reform wage as

$$w_B = \sum_{k \in B} \omega_k^N w_k, \quad \omega_k^N = \frac{\text{notifications from occupation } k}{\text{all B notifications}}. \quad (1.64)$$

This common w_B is used in every joint cell. The approximation is most credible when B occupations have similar wages and technologies and their notification mix is stable across worker types and time. Robustness exercises can restrict the comparison set or model its constituent occupations separately.

Difference-in-differences moments

The empirical DiD moments are the interaction coefficients from the paper's preferred specifications for applications and average applicant quality. The application coefficient comes from a Poisson model and is therefore a change in a log application rate. Its structural counterpart must use $h(x) = \log x$. The level and $\log(1+x)$ transformations are relevant only for alternative empirical specifications estimated in those units.

For quality, the empirical outcome remains the mean hiring score S_{ij} used in the reduced-form analysis. Its structural counterpart assigns $\chi_A(c)$ to an A application from cell c and $\chi_B(c)$ to a B application from that cell. Hence

$$\bar{S}_j = \frac{\sum_c \pi_j(c) \chi_j(c) \lambda_j(c)}{\sum_c \pi_j(c) \lambda_j(c)}.$$

The empirical and structural DiDs are therefore compared in score units. Percentile ranks $q_j(c)$ remain the inputs to the structural behavioral equations but are not substituted for the reduced-form dependent variable.

1.D.2 Formal identification

Let the directly observed or calibrated parameter vector be

$$\theta_0 = (\delta, \tilde{\beta}, \beta, M_A, M_B, \kappa_A, \kappa_B, \{\alpha_A(c), \alpha_B(c), \omega_c, q_A(c), q_B(c), \chi_A(c), \chi_B(c)\}_c),$$

and let

$$\theta = (b, b_1, c_A, c_B, s_A, s_B, \sigma_A, \sigma_B)$$

be the behavioural parameter vector.

We impose the following conditions. First, the time unit, mean match duration, wages, notification frequencies, notified workers per vacancy, and the empirical joint distribution of productivity ranks are observed. Second, the conditional moments are interior:

$$0 < \lambda_j(q) < 1, \quad 0 < p_j(q_j) < \kappa_j.$$

Each occupation contains at least two productivity values with distinct hiring probabilities. Third, the mean-field equilibrium is unique or a stable equilibrium-selection rule is fixed. Finally, the rank condition defined below holds.

Parameters identified directly

The baseline uses $T = 100$ model periods per day. This fine time grid is chosen so that

$$\alpha_A(c) + \alpha_B(c) \leq 1$$

in every empirical cell, as required by the mutually exclusive notification events in Section 1.3. Notification probabilities are constructed on this grid from times-tamped periods of unemployment and platform availability. If multiple notifications remain in a subperiod, the grid must be refined further rather than treating an arrival intensity above one as a probability.

Suppose mean match duration is L days. Geometric separation gives

$$\delta = \frac{1}{LT}. \quad (1.65)$$

Given annual discount rate r ,

$$\tilde{\beta} = \exp\left(-\frac{r}{365T}\right), \quad \beta = \frac{\tilde{\beta}}{1 - \tilde{\beta}(1 - \delta)}. \quad (1.66)$$

Wages and other flow payoffs are converted to the same $1/T$ -day unit before entering the Bellman equations. Thus increasing T changes the numerical time unit but not the economic present value of a job. Equation (1.22) identifies κ_j , while $\alpha_j(c)$,

1.D. STRUCTURAL ESTIMATION AND DECOMPOSITION

69

the cell weights, and M_j are direct notification moments linked by equation (1.63). If a Gaussian copula is imposed as a sensitivity check and τ_q is Kendall's rank correlation,

$$\rho = \sin\left(\frac{\pi}{2}\tau_q\right). \quad (1.67)$$

Employer-choice dispersion

Proposition 1.2 (Employer selectivity). *Conditional hiring probabilities $p_j(q_j)$, together with κ_j , point identify σ_j . For two distinct ranks q_j^1, q_j^2 ,*

$$\sigma_j = -\frac{q_j^2 - q_j^1}{\log[\kappa_j/p_j(q_j^2) - 1] - \log[\kappa_j/p_j(q_j^1) - 1]}. \quad (1.68)$$

Proof. The mean-field hiring equation implies

$$\log\left[\frac{\kappa_j}{p_j(q_j)} - 1\right] = \log D_j - \frac{q_j}{\sigma_j}.$$

The slope in productivity rank is $-1/\sigma_j$. Differencing between two ranks yields equation (1.68). $\square$

With more than two rank bins, linearity of transformed hiring odds and agreement of the intercept with the equilibrium competition index D_j provide overidentifying restrictions.

Worker-side parameters

Define the observable objects

$$G_j(q) = -\alpha_j(q)\beta p_j(q) \log[1 - \lambda_j(q)].$$

For each joint cell n , the two occupation equations can be written

$$w_A = b + b_1 \bar{q}_n + \frac{c_A}{p_{nA}} + s_A[\text{logit}(\lambda_{nA}) + G_{nA}] + s_B G_{nB}, \quad (1.69)$$

$$w_B = b + b_1 \bar{q}_n + \frac{c_B}{p_{nB}} + s_A G_{nA} + s_B[\text{logit}(\lambda_{nB}) + G_{nB}]. \quad (1.70)$$

Stack the equations as

$$W = X\theta_u, \quad \theta_u = (b, b_1, c_A, c_B, s_A, s_B)'$$

The A and B rows for cell n are

$$\begin{aligned} X_{nA} &= \left(1, \bar{q}_n, \frac{1}{p_{nA}}, 0, \text{logit}(\lambda_{nA}) + G_{nA}, G_{nB} \right), \\ X_{nB} &= \left(1, \bar{q}_n, 0, \frac{1}{p_{nB}}, G_{nA}, \text{logit}(\lambda_{nB}) + G_{nB} \right). \end{aligned}$$

Proposition 1.3 (Identification of worker-side parameters). *If X has column rank six, then*

$$\theta_u = (X'X)^{-1}X'W.$$

Thus the conditional application and hiring moments point identify $(b, b_1, c_A, c_B, s_A, s_B)$.

Proof. Equations (1.69)–(1.70) give $W = X\theta_u$. Full column rank implies that the null space of X is zero, so no two distinct parameter vectors produce the same moment collection. Premultiplying by $(X'X)^{-1}X'$ gives the inverse. $\square$

Theorem 1.1 (Mean-field identification). *Under the interiority, rank, observed-primitives, and equilibrium-selection conditions above, the mapping from admissible structural parameters (θ_0, θ) to the collection of aggregate, conditional application, and conditional hiring moments is one-to-one. Its inverse is given sequentially by equations (1.65)–(1.66), the directly observed notification and cell-distribution moments, Proposition 1.2, and Proposition 1.3.*

The DiD moments are not required for this constructive point-identification result. They overidentify the model and discipline its equilibrium response away from the pre-reform wage.

1.D.3 GMM moments

Let superscript M denote a model prediction at the pre-reform wage. For each joint cell, the standardised application moments are

$$g_{\lambda_{A,c}}(\theta) = \frac{\text{logit}(\hat{\lambda}_{A,c}) - \text{logit}(\lambda_{A,c}^M(\theta))}{\widehat{\text{se}}[\text{logit}(\hat{\lambda}_{A,c})]}, \quad (1.71)$$

$$g_{\lambda_{B,c}}(\theta) = \frac{\text{logit}(\hat{\lambda}_{B,c}) - \text{logit}(\lambda_{B,c}^M(\theta))}{\widehat{\text{se}}[\text{logit}(\hat{\lambda}_{B,c})]}. \quad (1.72)$$

For each distinct occupation-specific rank bin, the hiring moment is

$$g_{p_j,x}(\theta) = \frac{\log(\hat{\kappa}_j/\hat{p}_j(x) - 1) - \log(\hat{\kappa}_j/p_j^M(x;\theta) - 1)}{\widehat{\text{se}}[\log(\hat{\kappa}_j/\hat{p}_j(x) - 1)]}. \quad (1.73)$$

1.D. STRUCTURAL ESTIMATION AND DECOMPOSITION

71

The estimator also includes pre-reform applications per vacancy:

$$g_{\mathcal{A}_j}(\theta) = \frac{\hat{\mathcal{A}}_j - \mathcal{A}_j^M(\theta)}{\widehat{\text{se}}(\hat{\mathcal{A}}_j)}. \quad (1.74)$$

Let $\hat{\tau}_{\mathcal{A}}^{DID}$ and $\hat{\tau}_S^{DID}$ denote the empirical coefficients. The DiD moments are

$$g_{DID,\mathcal{A}}(\theta) = \frac{\hat{\tau}_{\mathcal{A}}^{DID} - \tau_{\mathcal{A}}^M(\theta)}{\widehat{\text{se}}(\hat{\tau}_{\mathcal{A}}^{DID})}, \quad (1.75)$$

$$g_{DID,S}(\theta) = \frac{\hat{\tau}_S^{DID} - \tau_S^M(\theta)}{\widehat{\text{se}}(\hat{\tau}_S^{DID})}. \quad (1.76)$$

Stacking all moments gives $g(\theta)$. The estimator minimizes

$$Q(\theta) = g(\theta)'Wg(\theta).$$

Positive parameters are optimised in logarithms. A constructive estimate based on Propositions 1.2 and 1.3 provides starting values; the numerical estimator uses several dispersed starts and retains the solution with the lowest criterion.

1.D.4 Weighting matrix and specification test

Let $\hat{\Omega}$ be the covariance matrix of the unstandardised empirical moments and let $\hat{S}$ be its diagonal matrix of standard deviations. Because g contains standardised residuals, the efficient weighting matrix is the inverse correlation matrix

$$W = \left(\hat{S}^{-1} \hat{\Omega} \hat{S}^{-1} \right)^{-1}.$$

We estimate $\hat{\Omega}$ from a common set of cluster-bootstrap draws that recompute worker ranks and score mappings when appropriate, notification and filling calibrations, cell probabilities, independently constructed aggregate moments, and both DiD estimates. All moments must carry the same bootstrap-draw identifier so their cross-covariances are estimable. When the empirical designs require more than one clustering dimension, the resampling scheme must reproduce that joint design or, equivalently, combine compatible cluster-level influence functions. A covariance matrix assembled from separate marginal standard errors is not sufficient.

At the optimum,

$$J = g(\hat{\theta})'Wg(\hat{\theta}) \stackrel{a}{\sim} \chi_{m-8}^2,$$

where m is the number of moments. We use its standard overidentification-test in-

terpretation only when the full joint covariance matrix has been estimated and the weighting matrix is nonsingular on the moment space. In addition to the asymptotic p -value, we obtain a finite-sample p -value from a recentered common cluster bootstrap under the fitted null, re-estimating all eight parameters and recomputing J in every draw. If these requirements are not met, J is reported only as a descriptive criterion value, not as a formal specification test.

1.D.5 Structural decomposition

For each outcome Y , solve the model at the pre- and post-reform wages:

$$\mathcal{E}^0 = \mathcal{E}(w_A^0, w_B; \hat{\theta}), \quad \mathcal{E}^1 = \mathcal{E}(w_A^1, w_B; \hat{\theta}).$$

These equilibria provide $\lambda_j^r(q)$, $p_j^r(q_j)$, and $u^r(q)$ for $r \in \{0, 1\}$. The occupation-level changes are

$$\Delta Y_j = Y_j(\mathcal{E}^1) - Y_j(\mathcal{E}^0),$$

and the structural DiD is exactly

$$\tau_Y^M = \Delta Y_A - \Delta Y_B.$$

When the empirical regression transforms the outcome, Y_j denotes that transformed outcome before differences are taken.

Channels within occupation B

At each point on the wage path, decompose the application response as

$$\frac{d\lambda_B^u(q)}{dw_A} = -\frac{\lambda_B(q)[1 - \lambda_B(q)]}{s_B} \frac{du(q)}{dw_A}, \quad (1.77)$$

$$\frac{d\lambda_B^p(q)}{dw_A} = \frac{\lambda_B(q)[1 - \lambda_B(q)]}{s_B} \frac{c_B}{p_B(q_B)^2} \frac{dp_B(q_B)}{dw_A}. \quad (1.78)$$

For applications per vacancy,

$$\frac{d\mathcal{A}_B^u}{dw_A} = M_B \sum_c \pi_B(c) \frac{d\lambda_B^u(c)}{dw_A}, \quad (1.79)$$

$$\frac{d\mathcal{A}_B^p}{dw_A} = M_B \sum_c \pi_B(c) \frac{d\lambda_B^p(c)}{dw_A}. \quad (1.80)$$

If the empirical outcome is $h(\mathcal{A}_B)$, each derivative is multiplied by $h'(\mathcal{A}_B)$.

1.D. STRUCTURAL ESTIMATION AND DECOMPOSITION

73

The average applicant score is

$$\bar{S}_B = \frac{\sum_c \pi_B(c) \chi_B(c) \lambda_B(c)}{\sum_c \pi_B(c) \lambda_B(c)}.$$

For a perturbation $d\lambda_B^k$, where $k \in \{u, p\}$, its directional derivative is

$$d\bar{S}_B^k = \frac{\sum_c \pi_B(c) (\chi_B(c) - \bar{S}_B) d\lambda_B^k(c)}{\sum_c \pi_B(c) \lambda_B(c)}. \quad (1.81)$$

The signs of these quality channels are unrestricted because they depend on which worker types change their application probabilities.

Integrating along the observed wage change gives

$$\Delta Y_B^k = \int_{w_A^0}^{w_A^1} \frac{dY_B^k(w)}{dw} dw, \quad k \in \{u, p\}. \quad (1.82)$$

The adding-up restrictions are

$$\Delta Y_B = \Delta Y_B^u + \Delta Y_B^p, \quad \tau_Y^M = \Delta Y_A - \Delta Y_B^u - \Delta Y_B^p.$$

1.D.6 Numerical solution and inference

For each trial parameter vector, the model is solved at w_A^0 . Worker application probabilities are obtained from the scalar continuation-value problem, conditional application probabilities are integrated using the joint productivity distribution, and hiring probabilities are updated from the mean-field competition indices. The iteration continues until the maximum change in log hiring probabilities is below the numerical tolerance.

The post-reform solution at w_A^1 is initialized from the pre-reform equilibrium. These two equilibria generate the model DiD moments. Numerical accuracy is assessed by tightening the productivity grid, fixed-point tolerance, and relaxation parameter.

For the exact-operator sensitivity exercise, we replace only the mean-field hiring update with the Poisson-logit operator in equation (1.43), evaluated by numerical quadrature. We then re-estimate all eight behavioral parameters rather than evaluating the exact operator at the mean-field estimate. The empirical moments, calibrated notification primitives, time unit, and wage counterfactual are identical across the two estimators.

For the channel decomposition, we solve the equilibrium on a fine grid between w_A^0 and w_A^1 . centred finite differences approximate du/dw_A and dp_B/dw_A , which are inserted into equations (1.77)–(1.81). We then use trapezoidal integration in

equation (1.82). The difference between the total B response and the sum of the two channels is reported as a numerical adding-up error.

For inference, each bootstrap draw repeats the complete procedure:

1. resample using the common clustering design for all empirical objects;
2. reconstruct the notification counts, ranks, score mappings, pre-reform moments, aggregate application moments, and DiD estimates;
3. recalibrate the directly identified parameters;
4. re-estimate the eight behavioural parameters;
5. solve the pre- and post-reform equilibria;
6. compute model fit, J , and the structural decomposition.

Percentile or studentized intervals from this distribution incorporate sampling uncertainty in both the reduced-form and structural stages.

Chapter 2

Caseworkers' Meetings with Job-seekers: Meeting Timing and the Gender Gap in Returns to Counselling

Abstract

Public employment services rely on mandatory meetings between caseworkers and unemployed job-seekers. A large literature examines their frequency and content, but what shapes the quality of a given meeting is harder to observe. We ask whether the time of day at which a meeting takes place relates to job-seekers' outcomes, motivated by evidence that time of day affects the quality of evaluative decisions in other settings, and we examine whether any such association differs across job-seekers.

We use Swiss unemployment insurance records from SECO covering all individuals registered as unemployed between 2001 and 2018. After cleaning, the sample comprises more than 11 million meetings, each recorded with an exact timestamp, involving roughly 1.2 million job-seekers and 6,300 caseworkers. Meeting times are set by caseworkers and accepted by job-seekers rather than self-selected. We compare outcomes across meetings held at different hours by the same job-seeker within a single unemployment spell, using individual-by-spell fixed effects. Early-morning meetings are associated with a coherent pattern across outcomes: higher employment six months later, faster exit from unemployment insurance, lower benefit receipt, and reduced participation in labour market programmes. This pattern suggests faster placement into employment rather than changes in job quality. Examining **heterogeneity, we find that the association is more pronounced** for women, and that it is specific to

This chapter is co-authored with Conny Wunsch (University of Basel).

the early-morning slot. The magnitudes are modest, but the consistency of the pattern across outcomes suggests that appointment scheduling, a routine and essentially costless administrative choice, may affect job-seekers unequally.

2.1 Introduction

Public employment services deliver job search assistance largely through one instrument: mandatory meetings between caseworkers and unemployed job-seekers. A substantial literature establishes that these meetings are a productive input, and identifies who should conduct them and how often: outcomes respond to caseworker quality and counselling style, to the frequency and sequencing of meetings, and to caseloads (Behncke et al., 2010b; Schiprowski, 2020; Van den Berg et al., 2012; Maibom et al., 2017; Böheim et al., 2022). Yet the meeting itself remains a black box. Administrative registers record that an interaction took place, not what happened inside it, and so we know little about what makes a meeting effective.

One possible natural window into that black box is meeting duration, the time a caseworker devotes to a job-seeker, and this is indeed where our project started. As Section 2.3 documents, however, the recorded durations cannot support such an analysis: they are reliably recorded but mostly concentrated in a few values. The register does observe one dimension of the interaction that the literature has never used: the hour at which the meeting is scheduled. Evidence from other high-stakes settings suggests this lever should not be presumed neutral: students score higher in morning exams (Pope, 2016), and recruiters on a Swiss public employment platform screen applications less carefully before lunch and at the end of the day, amplifying discrimination against minority candidates (Hangartner et al., 2021).

The question this chapter asks is therefore broader than whether early meetings beat late ones. It is whether apparently trivial choices in the operational design of a public service, such as the distribution of appointments over the working day, can change both the effectiveness of an active labour market policy and the distribution of its benefits. We first ask whether the time of day at which a caseworker meeting takes place affect the subsequent labour market outcomes of the unemployed. Then we investigate whether such effect is distributed evenly across workers.

We answer these questions using Swiss unemployment insurance registers covering all registered unemployed between 2001 and 2018: more than 11 million meetings with exact timestamps, linked to employment, earnings, and benefit outcomes at fixed horizons after each meeting. Identification exploits within-individual, within-spell variation in meeting hour. Our preferred specification includes individual-by-spell fixed effects and caseworker fixed effects, so that the comparison is between meetings of the same job-seeker, within the same unemployment spell, at the same stage of the spell, held at different times of day, net of any caseworker-level differences in scheduling and effectiveness. The residual variation in timing plausibly reflects caseworker-side scheduling constraints (agenda

availability, substitutions, office reorganisation) rather than job-seeker choice, and we probe this interpretation with a battery of placebo tests on meetings that were scheduled but never took place.

Meeting timing matters, but modestly: an early-morning (7–8 a.m.) meeting is associated with a 0.22 percentage point higher probability of employment six months later and 0.019 additional months of employment over the following half-year, relative to the 9–10 a.m. reference slot. This modest average conceals substantial heterogeneity. The early-morning advantage is roughly twice as large for women as for men (+0.32 vs. +0.15 percentage points on six-month employment; interaction $p = 0.021$), and the gap is larger and more precisely estimated on cumulative employment. It is, moreover, *specific* to the early-morning slot: interactions between gender and every other time slot are precise zeros. Women do not simply do better at all meetings; the advantage is concentrated in the early-morning slot alone. This specificity makes a data-mining explanation less likely, since a chance rejection would not single out one slot, and the pattern is stable across alternative clustering, the exclusion of the rare evening slot, and an alternative outcome, as well as within the subsample of job-seekers whose spells contain genuine within-spell variation in meeting timing.

Our contribution is to open one dimension of this black box: intra-day scheduling, a margin of policy delivery that is observable, administratively set, and cheap to change, yet that the literature has largely overlooked. The average return there is modest; what stands out is its distribution. An administrative choice with no distributive intent concentrates the benefits of the most effective slot among women, so the operational design of the service shapes not only whether the policy works but for whom.

We read the gender result as an equity finding rather than as evidence of discrimination, and the estimates as reduced-form, since the data do not let us identify the mechanism behind the timing effect; we return to both points in Section 2.8.

The rest of the chapter is organised as follows. After a review of the literature 2.2, Section 2.3 sets out the institutional context and the register data, and Section 2.4 the empirical strategy, together with the identification threats, dynamic selection foremost among them. The average effects come first, in Section 2.5, before Section 2.6 turns to the gender heterogeneity. Section 2.7 gathers the robustness and placebo evidence, including the single- versus multi-spell decomposition, and Section 2.8 draws out the interpretation and the policy implications. Section 3.6 concludes.

2.2. LITERATURE REVIEW

2.2 Literature review

The interaction between job-seekers and caseworkers is by now well documented to affect unemployment duration and later labour-market outcomes, through several distinct channels that this section reviews. As caseworkers have the dual responsibility of counseling and monitoring unemployed persons, it is relevant which type of caseworker is assigned to the unemployed person.

Behncke et al., 2010a find that similarity between the caseworker and the unemployed person along several socio-economic dimensions improves employment prospects. No single dimension, whether gender, age, nationality or educational background, is sufficient on its own; it is similarity across several characteristics jointly that matters. Arni, 2015 highlights the difficulty of disentangling the counselor's non-cognitive, motivational and psychological role from the effect of the activation measures themselves, since the two operate jointly.

The choice of the type of monitoring has also a non negligible relevance. On the one hand too strict caseworkers, applying "sticks" treatment rather than "carrots", results in an ineffective performance and reduced earnings (Arni et al., 2020). On the other hand, less cooperative counselors have greater placement rates in Switzerland Behncke et al., 2010b, probably due to factors such as threats of fines and pressure to take jobs Huber et al., 2017. Berg and Klaauw, 2006 suggest more stringent monitoring leads to more formal job search but less informal job search, with a mixed effect on unemployment length and job entrance rates.¹

The value of in-person interaction has been examined by Schiprowski, 2020. Accounting for the incidence of unplanned caseworker absence, she has measured the impact of the face-to-face meetings. Furthermore, she documents heterogeneity within caseworkers' quality. For lower-than-median productivity caseworkers, the absence's effect is zero. On the contrary, when a job-seeker involuntary loses a meeting with a caseworker in the upper half of the distribution, the unemployment duration increases. More recent work relates this face-to-face channel to job-search behaviour itself: Schiprowski et al., 2024 show that caseworkers affect the search effort exerted by job seekers, pointing to search intensity as one mechanism through which meetings operate.

Following this line of research, Dromundo Mokrani, 2021 estimate the value added by caseworkers in terms of getting a job quickly and more stable employment. They find that fast job placement success does not always mean the placing of job seekers in top-quality positions. Exploiting the date-of-birth assignment rule used by many Swedish employment offices as a source of as-if random variation, Cederlöf et al., 2025 document substantial caseworker value-added on both

¹For a more detailed discussion see e.g. Arni and Schiprowski, 2019 and Manning, 2009

the speed of re-employment and the quality (wages) of the subsequent job, and find that more experienced and more "active" caseworkers, those who meet their clients more frequently, obtain better outcomes.

This last point, that caseworkers who meet their clients more frequently obtain better outcomes, connects to a broader question: the meeting's effect is not just driven by caseworkers but also depends on the duration as well as the frequency.

The transition rate from unemployment to employment strongly increases in the week after the meeting and this effect persists for some weeks after, but eventually fades. After the second meeting, the transition rate increased again and even more than after the first time. For the third meeting, it peaks at an even higher level, and so on for every meeting. The size effect on the finding job rate increases with the number of meetings (Van den Berg et al., 2012).

Maibom et al., 2017 find that the frequent and early meetings between newly unemployed workers and caseworkers increase the employment rates over the next two years by up to 5 weeks, corresponding to a 10% employment effect. The findings are heterogeneous by gender. The effect of participating in early active labor market programs has a negative effect for men, but not on women. Furthermore, some nudging effects have been registered based on the timing of meetings to sign so-called Integration Agreements (Van den Berg et al., 2020). Whether meetings are effective even before entry into unemployment is less clear. In a field experiment offering preventive meetings to German workers who were still employed but at risk of job loss, Homrighausen and Oberfichtner, 2025 find that the intervention brings the first meeting forward and increases the total number of meetings, but does not improve unemployment or earnings outcomes, and does not alter job-search behaviour.

Fewer caseloads in public employment offices have been observed to be associated with an increase in meetings between jobless people and their caseworkers, job offers, program assignments, and sanctions for failing to follow job search requirements (Böheim et al., 2022).

In the empirical literature, it has been found that the time of day affects the productivity of job recruiters Hangartner et al., 2021, who spend less time on applications before lunch and towards the end of the day, following a cycle related to fatigue. Beyond this setting, a large literature documents time-of-day effects in other domains, most prominently in education and law. Pope, 2016 examines how the time of day impacts students' productivity, and the results suggest that pupils learn more in the morning than they do later in the day, and consequently, their test scores are higher in the morning as opposed to the afternoon. Haggag et al., 2021 finds that students enrolled in the early-morning session of a course are less likely to continue their studies in that field and major in the same area. Gaggero

2.3. DATA

81

and Tommasi, 2023 finds a hump shape relationship between exam scores and timing. In the same vein, Hervé et al., 2024 show that being tested shortly after a meal reduces adolescents' cognitive performance, consistent with a post-prandial dip.

A related strand documents decision fatigue in high-stakes settings. Danziger et al., 2011 find that the share of favourable parole rulings declines over the course of a session and rebounds after meal breaks, which they read as evidence of decision fatigue. However, Shroff and Vamvourellis, 2022, re-examining nearly 100,000 pretrial-release decisions, attribute most of this within-day decline to the changing composition of cases over the day and find no rebound after breaks, casting doubt on the generality of the fatigue interpretation.

These candidate mechanisms need not push in the same direction. Decision fatigue and post-meal dips predict deterioration as the working day, or the interval since a meal, progresses, whereas circadian factors can instead penalise performance in the early hours: Heissel and Norris, 2018 show that later school start times relative to sunrise raise adolescents' test scores, an effect they attribute to circadian misalignment. Whether time of day shapes the outcomes of caseworker meetings, and whether it does so unequally across job-seekers, has to our knowledge not been examined, and this is the question we address.

2.3 Data

2.3.1 Institutional setting

Job seekers registered with the Swiss public employment service are legally required to attend regular counseling meetings at their local employment office (RAV). The first meeting takes place within two weeks of registration, and meetings are held at least every two months thereafter. For our design, the dates are set by the employment office rather than chosen by the job seeker, and unexcused absences are sanctioned through temporary benefit suspensions.² This administrative control over timing is what the identification strategy exploits; Section 2.4 returns to the limited margin the job seeker retains.

The empirical analysis is based on the Swiss unemployment insurance (UI) administrative register provided by the State Secretariat for Economic Affairs (SECO), which covers all persons officially registered as unemployed between 2001 and 2018 and is built from two linked information systems, the AVAM (placement

²The relevant provisions are contained in the Unemployment Insurance Act (AVIG; SR 837.0) and the accompanying ordinance (AVIV; SR 837.02); Appendix 2.A sets out the detailed legal framework, including the provisions governing the scheduling and rescheduling of meetings. The institutional setting is also described in Schiprowski (2020).

and labour market statistics) and the ASAL (benefit payments), with individuals matched across systems via social security number. The full register contains 20,425,634 meetings between 1,259,358 job-seekers and 7,397 caseworkers, each with an exact timestamp and a recorded duration, alongside 247 variables covering job-seeker sociodemographic characteristics, employment history, caseworker characteristics, meeting type, and meeting timing (Appendix Tables A8 and A9 describe the variables). We restrict this register to an analysis sample of 11,745,735 meetings in two steps, detailed in Table A7: we keep only meetings that took place, recorded at plausible hours (7 a.m. to 7 p.m.), and then those with plausible recorded durations (20–120 minutes, which removes the pro-forma entries massed at the 0–20 minute administrative default). Meetings that did not take place (postponed, cancelled, or not attended) are set aside for the placebo analysis of Section 2.7.2, and Appendix Tables A14–A15 document the implausible-hour and implausible-duration records. The fixed-effects specifications further drop observations that become singletons once the fixed effects are absorbed, leaving 9,025,553 meetings under individual fixed effects and 8,832,592 once individual-by-spell fixed effects are absorbed, since a spell must contain more than one meeting at distinct hours to contribute within-spell variation.




Figure 2.1: Distribution of meetings over the working day, by meeting status
 Notes: Share of meetings (in %) held in each one-hour slot from 7 a.m. to 6 p.m., shown separately for each recorded meeting status. The 6 p.m. slot appears empty but is not.

The treatment is the two-hour time slot in which the meeting takes place, aggregated to two-hour windows because the results are more homogeneous and more precisely estimated at that level: 7–8 a.m. (the focal early-morning slot), 9–10 a.m. (the most frequent slot and the reference category throughout), 11–12 a.m., 1–2 p.m., 3–4 p.m., and the rare and noisy 5–6 p.m. slot (roughly 0.4% of meetings), with day-of-week dummies included throughout and Tuesday as reference. Figure 2.1 shows how meetings are distributed over the working day: the 9–10 a.m.

slot is the most frequent, the early-morning and late-afternoon slots are progressively rarer, and the profile is almost identical across meeting statuses.

A second treatment dimension would have been meeting duration, which the register records for every meeting and which we initially set out to study alongside timing, since the time a caseworker devotes to a job-seeker is arguably the channel through which differential treatment across groups would operate most directly. After cleaning, however, the recorded durations are concentrated at a small number of mass points, with pronounced spikes at exactly 30, 45, and 60 minutes: they are reliably recorded, but only at a few rounded values rather than as a precise, continuous measure. This is not a data-quality problem, but it leaves too little variation within each duration group to identify a marginal effect of meeting length on outcomes.

We therefore treat duration as a reliable categorical variable rather than a continuous one, using it to trim implausible records and entering the duration category as a control in the regressions, and we return to the intensity margin in the discussion rather than reading these categories as a measure of meeting effort.

Outcomes are constructed at the meeting level at fixed horizons after the meeting date. The main outcome, employment at six months, equals one if the job-seeker is employed six months after the meeting, while the second main outcome, cumulative months employed over the following half-year, sums six binary monthly employment indicators and so ranges from 0 to 6; the first is a snapshot and the second captures how quickly the job-seeker left unemployment, so reporting both lets us see whether the effect of morning meetings works through faster exit, through higher employment at a fixed horizon, or both. The remaining outcomes follow the same construction: cumulative months registered on UI, cumulative UI benefits received in CHF, monthly earnings at month six, cumulative months in active labour market programmes, and cumulative months with interim earnings.

The six-month horizon of these outcomes is distant from the meeting itself. We are extending the analysis to monthly horizons using a newly assembled version of the register data; this remains ongoing work, no estimate from it enters the present chapter, and we return to it in the conclusion.

2.4 Empirical strategy

2.4.1 Fixed-effects specifications

Our analysis identifies the effect first within job-seekers, then within unemployment spells, and finally conditioning on caseworkers. As a preliminary and de-

scriptive step, Appendix 2.B reports a baseline OLS specification with sequentially added controls and caseworker fixed effects. That exercise estimates a different quantity from the one estimated here: it compares across job-seekers rather than within them, and its outcome is anchored to the unemployment spell rather than to the meeting. The raw association between the 7–8 a.m. slot and employment is small and insignificant; once job-seeker characteristics are added, it is positive and stable at 0.29 to 0.38 percentage points, and it is this regularity that motivated the design set out below. We draw no causal conclusion from it. Coefficient stability with respect to observables is not evidence of stability with respect to unobservables (Oster, 2019), and the scheduling decision rests with the caseworker: if she systematically assigns morning slots to job-seekers she perceives as more motivated, the coefficient of interest would capture not the effect of meeting timing but the effect of these unobservable characteristics. Nothing in that specification rules this out, since it compares different job-seekers rather than the same job-seeker with herself. The OLS therefore serves only to establish the descriptive pattern; the fixed-effects strategy addresses this by comparing a job-seeker with herself rather than with others, and removes successive layers of potential confounding. We estimate three nested specifications, pooling all meetings and absorbing the month of the spell as a fixed effect.³

The unit of observation in the analysis is a meeting. We keep every meeting that took place at a plausible hour and with a plausible recorded duration, as described in Section 2.3. A job-seeker therefore enters the sample once for each meeting she attends: a job-seeker who attends five meetings contributes five observations. Spells are not aggregated either. A job-seeker who registers as unemployed more than once over the period we observe enters the sample with the meetings of all her spells, and not only with those of her first one. We index meetings by m and job-seekers by i , and for a meeting m attended by job-seeker i we write $s(i, m)$ for the spell during which it takes place, $c(i, m)$ for the caseworker who conducted it, who is not necessarily the one to whom i was assigned at registration, and $t(i, m)$ for the month of that spell in which the meeting falls, counted from the start of the spell and therefore restarting at each new spell of the same job-seeker.

Meetings and months are not in one-to-one correspondence, and this is worth stating because the specifications below absorb $t(i, m)$ as a fixed effect. Months of a spell in which no meeting takes place generate no observation: meetings are approximately, but not exactly, monthly, and a meeting-free month does not enter the estimation sample. Conversely, a month may contain more than one meeting, in which case each of them is a separate observation sharing the same value of

³Estimation relies on the high-dimensional fixed-effects estimator of Correia and Constantine (2026).

$t(i, m)$. A job-seeker with two meetings contributes two observations whether or not they fall in the same month of the spell.

The outcome is measured six months after the meeting. Y_{im}^6 equals one if job-seeker i is employed six months after meeting m , the superscript recording the horizon at which the outcome is measured.⁴ The second main outcome counts the months of employment recorded over the six months that follow meeting m and therefore ranges from zero to six: it registers how much employment follows the meeting, while the former registers only whether the job-seeker is employed at the end of the window. The remaining outcomes described in Section 2.3 are of one of these two kinds, a status at the end of the window or a count over it, and enter the same specifications in turn.

Version 1: job-seeker fixed effects.

$$Y_{im}^6 = \beta_1 \text{time}_{im} + \beta_2 \text{day}_{im} + \beta_3 Z_{im} + \beta_4 W_{c(i,m)} + \beta_5 X_{i,s(i,m)} + \alpha_i + \lambda_{t(i,m)} + \gamma + \varepsilon_{im} \quad (2.1)$$

The coefficient on the 7–8 a.m. slot is the object of interest and the estimate discussed in the text; Section 2.5.2 reports the coefficients on the remaining slots. time_{im} denotes the two-hour slot in which meeting m takes place, as defined in Section 2.3, with the 9–10 a.m. slot as reference category, and day_{im} the day of the week, with Tuesday as reference. Z_{im} collects the characteristics of the meeting entered alongside the treatment: its recorded duration category, its type, and an indicator for whether it was conducted by the caseworker to whom i was assigned at registration rather than by a substitute standing in for her, since the quality of the interaction may differ when the two already know each other. $W_{c(i,m)}$ collects characteristics of the caseworker conducting the meeting: average caseload, work experience, the average duration of her meetings, and the share of her meetings with missing recorded duration.

$X_{i,s(i,m)}$ contains the job-seeker's sociodemographic characteristics and employment history, measured at the start of the spell: gender, age, education, nationality, language skills, marital status, geographic mobility, the workload and hierarchical position of the last job, previous earnings and tenure, receipt of child and training allowances and of disability benefits, benefit entitlement, the share of time in dependent employment before the current spell (a proxy for pre-unemployment labour market attachment), indicators for any prior self-employment, disability, unemployment, or other non-employment spells, and detailed employment, regis-

⁴We index the outcome by the meeting rather than by the month of the spell, and we do so deliberately: writing it as Y_{it} would invite the reader to understand it as the status of job-seeker i in month t , whereas it is her status six months later.

tration, and benefit histories over the preceding 24 and 60 months. Because these characteristics are fixed within a spell, they are identified here only through variation across the spells of the same job-seeker, and they drop out of Version 2 and Version 3 below.

Finally, α_i is a job-seeker fixed effect and $\lambda_{t(i,m)}$ a fixed effect for the month of the spell, so that meetings are compared at the same stage of the spell, while γ collects fixed effects for economic sector, occupation group, canton, and the calendar month of the year in which the meeting takes place. Identification comes from within-job-seeker variation in meeting timing, both within and across unemployment spells, net of the stage of the spell; relative to Version 2 below, the across-spell comparisons are the additional source of variation that this specification uses. Meetings that become singletons once the fixed effects are absorbed drop out of the estimation, which is why the number of observations differs across the three specifications. Standard errors are clustered at the job-seeker level. Variable definitions are reported in Appendix Tables A8 and A9.

Version 2: job-seeker $\times$ spell fixed effects.

$$Y_{im}^6 = \beta_1 \text{time}_{im} + \beta_2 \text{day}_{im} + \beta_3 Z_{im} + \beta_4 W_{c(i,m)} + \alpha_{i,s(i,m)} + \lambda_{t(i,m)} + \gamma + \varepsilon_{im} \quad (2.2)$$

Version 2 replaces the job-seeker fixed effect α_i with a job-seeker $\times$ spell fixed effect $\alpha_{i,s(i,m)}$; all other terms in equation (2.1) are unchanged, except that the job-seeker characteristics $X_{i,s(i,m)}$, measured at the start of the spell and thus collinear with $\alpha_{i,s(i,m)}$, drop from the estimation. Identification now comes from comparing meetings of the same job-seeker, in the same unemployment spell, at the same month of the spell, but at different times of day: anything that is constant within a spell, such as her motivation on entering unemployment or the difficulty of her case, is absorbed by $\alpha_{i,s(i,m)}$, while $\lambda_{t(i,m)}$ absorbs the average change in reemployment prospects over the course of the spell. A spell contributes to the estimation only if it contains at least two meetings held at distinct hours.

Version 3: adding caseworker fixed effects (preferred).

$$Y_{im}^6 = \beta_1 \text{time}_{im} + \beta_2 \text{day}_{im} + \beta_3 Z_{im} + \beta_4 W_{c(i,m)} + \alpha_{i,s(i,m)} + \mu_{a(i,m)} + \lambda_{t(i,m)} + \gamma + \varepsilon_{im} \quad (2.3)$$

Version 3 adds fixed effects $\mu_{a(i,m)}$ for the caseworker responsible for the job-seeker, on top of the job-seeker $\times$ spell fixed effects. These absorb any time-invariant feature of that caseworker, including the composition of her caseload and her scheduling habits: if certain caseworkers both rely more on morning slots and

achieve better outcomes on average, that is captured by $\mu_{a(i,m)}$. A descriptive check in Section 2.4.2 complements this: caseworkers who rely more on the early slot do not handle observably stronger job-seekers. We condition on the responsible caseworker rather than on the one who actually conducted the meeting, since responsibility is settled before any individual meeting is scheduled, whereas who ends up attending a given meeting is decided meeting by meeting and could itself respond to the slot; the two coincide for 95% of meetings, the remainder being those held with a substitute.

Version 3 is our preferred specification and the headline estimates in the text refer to it, since it is the most demanding of the three, jointly conditioning on unobservable job-seeker and caseworker heterogeneity. That the coefficient on the 7–8 a.m. slot barely moves between Versions 2 and 3 (from 0.0023 to 0.0022 on employment six months after the meeting) is itself informative, indicating that the within-spell result of Version 2 is not driven by caseworker-level sorting into time slots. Both versions are reported throughout. Estimation is on the full sample; some robustness exercises, namely the spell split of Section 2.7.1 and the placebo tests of Section 2.7.2, are estimated on Version 2, and given the near-identity of the two versions the results would be unchanged under Version 3. Section 2.7.1 shows that restricting to job-seekers with at least two unemployment spells changes the point-in-time employment coefficient for reasons of sample composition rather than identification, and that cumulative outcomes are stable across samples.

2.4.2 Identification and threats to validity

Our empirical strategy rests on the assumption that the timing of caseworker meetings contains variation that is plausibly exogenous to unobservable job-seeker characteristics. We cannot claim random assignment of meeting times. Three distinct concerns deserve separate treatment. Selection into scheduling refers to the possibility that caseworkers systematically assign job-seekers with better unobservable characteristics (higher motivation, stronger labour market attachment, greater discipline) to morning slots, in which case the morning coefficient would capture the characteristics of those assigned rather than the effect of timing itself. Selection into realization is a different concern: conditional on being scheduled for a morning meeting, job-seekers who actually attend may differ systematically from those who do not, so that the sample of realised morning meetings would be positively selected. Dynamic selection over the spell, finally, arises because meetings at later spell stages are observed only for job-seekers who remain unemployed up to that point.

The institutional setting bounds the scope for this concern. In the Swiss pub-

lic employment service the meeting time is set by the caseworker, who takes the job-seeker's availability into account but does not delegate the choice to her: the job-seeker's influence on the slot operates through availability constraints (work schedules, care obligations, commuting), which are largely observable and included in our controls.

The data show no trace of sorting on observables, on either side of the scheduling decision. On the job-seeker side, the distribution of all observable characteristics (employment history, earnings, education, age, nationality, family situation) is balanced across time slots (Appendix Tables A10–A13). On the caseworker side, the share of a caseworker's meetings scheduled at 7–8 a.m. is uncorrelated with the average observable strength of her job-seekers (correlation with mean previous earnings -0.003 , with mean earnings over the past 60 months $+0.024$), and the mean previous earnings of their job-seekers are flat across quintiles of caseworkers' reliance on the early slot.⁵ These are descriptive checks on observables, and they do not speak to selection on unobservables; two further pieces of evidence do. First, Version 3 absorbs caseworker fixed effects on top of individual $\times$ spell fixed effects: if specific caseworkers systematically assigned morning slots to stronger job-seekers and also achieved better outcomes, this would be absorbed by μ_c , yet the 7–8 a.m. coefficient is virtually unchanged between Versions 2 and 3 (0.0023 to 0.0022). Second, were caseworkers sorting stronger job-seekers into morning slots, those slots should predict outcomes even for meetings that never took place, since those job-seekers would find employment regardless. We find no significant effect of the 7–8 a.m. slot for canceled or not-attended meetings. These non-realised samples are much smaller than the main one, so the placebos lack the power to separate a zero from an effect the size of our estimate, and they may also disproportionately represent weaker job-seekers, which could mechanically attenuate the coefficient; we therefore read them as consistent with the absence of such sorting rather than as definitive. The identifying weight rests instead on the decomposition of the postponed placebo by source of postponement, which locates the residual selection in the postponement decision rather than in scheduling (full results in Section 2.7.2).

Dynamic selection over the spell is the third concern, and the one our design addresses least directly. Because 87% of meetings are the only meeting of their job-

⁵Quintiles here rank the 4,746 caseworkers with at least 50 meetings in the analysis sample by the share of their meetings scheduled at 7–8 a.m., from the bottom quintile (which rarely uses the early slot) to the top (which uses it most). Mean previous earnings of their job-seekers are similar at the two ends of this distribution (4,579 CHF in the bottom quintile, 4,581 in the top). Previous earnings are only a partial summary of a job-seeker's observable strength, and we do not rely on this comparison alone: caseworker-side sorting is chiefly addressed by the caseworker fixed effects in Version 3, which absorb any time-invariant feature of a caseworker including the composition of her caseload.

seeker-spell in a given month of the spell, the within-spell variation that identifies the effect comes largely from comparing the same job-seeker at different stages of the spell rather than from meetings held side by side, and later stages are observed only for those who have not yet exited, a group selected toward the harder to place. The month-of-unemployment fixed effect λ_t removes the average outcome gradient over the spell, but not a residual correlation between which meeting carries the deviating hour and the outcome. A first version of the concern is bounded by the design itself: dynamic selection can contaminate the 7–8 a.m. coefficient only if the early slot falls at systematically different spell stages than the reference, whereas when 7–8 a.m. and 9–10 a.m. meetings occur at similar stages, survivor composition applies equally to treatment and reference and cancels in the within-spell contrast, and the flat balance of observable histories across slots (Tables A10–A13) is consistent with such a common stage distribution. Two residual concerns remain: the effect of meeting timing itself may vary over the spell and the motivation may fluctuate within individuals over time. Neither possibility is absorbed by the additive fixed effects. We treat these as unlikely drivers rather than ones we have logically excluded, and lean on a single feature that speaks against all three: any such artefact would have to concentrate in the 7–8 a.m. slot, vanish at every other hour, and double for women (Section 2.6), a conjunction none of them has reason to satisfy. The clean confirmation, which we leave to future work, is to re-estimate the early-morning effect within the first months of the spell, where little exit has yet occurred.

2.5 Main results: average effects of meeting timing

2.5.1 Average effects on the full sample

Table 2.1 reports the effect of a 7–8 a.m. meeting on all outcomes, estimated on the full sample using the three fixed effect specifications of Section 2.4. The results are consistent across outcomes and highly stable between columns (2) and (3), supporting that the estimates are not driven by caseworker-level sorting into morning slots.

Morning meetings increase the probability of employment at six months by 0.22 to 0.38 percentage points and the cumulative months of employment in the six months following the meeting by 0.019 to 0.028 months, with both coefficients significant at the 1% level in every specification. The reduction in magnitude from column (1) to column (2) reflects the elimination of between-individual selection: once individual fixed effects are replaced by individual $\times$ spell fixed effects, only within-spell variation in meeting timing identifies the effect, which is more conser-

Table 2.1: Effect of a 7–8 a.m. meeting on all outcomes (full sample)

Outcome	(1) Individual FE	(2) Indiv.×Spell FE	(3) Indiv.×Spell + CW FE
<i>Employment</i>			
Employed at month 6	0.0038*** (0.000432)	0.0023*** (0.000370)	0.0022*** (0.000369)
Cum. months employed	0.0275*** (0.00191)	0.0195*** (0.00153)	0.0191*** (0.00153)
<i>Unemployment insurance</i>			
Cum. months on UI	−0.0200*** (0.00189)	−0.0107*** (0.00154)	−0.0102*** (0.00153)
Cum. benefits paid	−90.48*** (6.629)	−40.69*** (5.479)	−38.83*** (5.451)
<i>Income</i>			
Income at month 6	14.88*** (2.337)	1.283 (2.016)	0.726 (2.009)
<i>Programme participation</i>			
Cum. months in ALMP	−0.0053*** (0.00149)	−0.0197*** (0.00138)	−0.0198*** (0.00137)
Cum. months, interim earn.	0.0046*** (0.00154)	−0.0101*** (0.00131)	−0.0102*** (0.00130)
Individual FE	✓		
Individual×Spell FE		✓	✓
Caseworker FE			✓
Month of unempl. FE	✓	✓	✓
Additional FE	✓	✓	✓
Clustered SE	✓	✓	✓
Observations	9,025,553	8,832,592	8,832,292

Notes: Standard errors clustered at the job-seeker level in parentheses. The coefficient reported is that on the 7–8 a.m. dummy, relative to the reference slot 9–10 a.m. All outcomes are measured over the six months following the meeting; cumulative outcomes sum months or amounts over that window. “Employed at month 6” is a 0/1 indicator; benefits paid and income are in CHF; all other outcomes are in months. Additional FE include occupation group, canton, and calendar month of the meeting. The reduction in observations between column (1) and columns (2)–(3) reflects the absorption of individual×spell fixed effects, which requires within-spell variation in meeting timing. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

2.5. MAIN RESULTS: AVERAGE EFFECTS OF MEETING TIMING

91

vative. The negligible change between columns (2) and (3) confirms that the result is not driven by caseworkers who systematically assign morning slots to more employable job-seekers. The unemployment insurance outcomes mirror this pattern: morning meetings reduce cumulative months on UI by 0.010 to 0.020 months and cumulative UI benefits received by CHF 39 to 90, consistent with an effect that operates through earlier exit from unemployment.

Income here is measured over all months after the meeting, counting zeros for months out of work, so it reflects both whether people work and how much they earn when they do. The coefficient is positive under individual fixed effects (column 1) but falls close to zero within spells (0.73 CHF in column 3), so the column-1 figure was between-individual selection, not an effect of timing. The near-zero is meaningful because of what it is missing. The employment gain of 0.0022 should by itself push unconditional income up by about 13 CHF, simply because more people are working (0.0022 times the average wage of the employed, about 5,900 CHF). Yet the estimate is 0.73 CHF. If more people work but income does not rise even by that much, the extra jobs are likely low-paid ones⁶. Taken together, the pattern suggests faster placement without income gains, or with no better job quality.

The within-spell specifications also show reductions in cumulative participation in active labour market programmes and interim employment; as discussed in Section 2.8, these cumulative outcomes are measured over months at risk and therefore mechanically decline when exit from unemployment is accelerated, so we do not interpret them as separate channels. The average effect is modest, and we present it as such. A 0.019-month increase on a base of 1.56 months (about 1.2% of the mean, or roughly half an extra day of employment over the half-year), and a 0.22 percentage point increase on a mean of 0.38 (about 0.6% of the mean), are small in absolute terms. That base of 0.38 is itself a stock: employment at six months is measured at a fixed horizon after each meeting and pooled over all meeting-months of the spell, so it is the share employed at that horizon, not a final reemployment rate, and the percentages of the mean quoted here should be read against it. Small as it is, the effect deserves a closer look. Reallocating appointment times has no budgetary cost, so even a small effect implies a favourable cost-benefit ratio. Added up over everyone, this effect could look large, and even attractive next to expensive active labour market policies that buy only small gains. But it is not free to harvest: not everyone can be moved into a morning slot, so the aggregate cannot simply be scaled up. We return to this in Section 2.8. Table 2.1 therefore does not say “no effect”; it shows a modest average that masks

⁶We say “likely” rather than “clearly” because income is very noisy (standard deviation 2,871 CHF) and the estimate is imprecise, so we cannot rule out a small positive effect.

substantial heterogeneity.

2.5.2 Effects across all time slots

Table 2.2 reports coefficients for all time slots relative to the 9-10 a.m. reference. The pattern across slots is informative both substantively and as a further check on the identification strategy.

The 11-12 a.m. slot produces a near-zero and insignificant coefficient on employment at six months, but the coefficients on cumulative months of employment and on cumulative UI registration remain significant and in the expected direction, suggesting that midday meetings generate some acceleration in exit from unemployment that does not persist to a detectable employment advantage at the fixed six-month horizon. The 1-2 p.m. slot shows a precisely estimated null on six-month employment and a near-zero coefficient on cumulative employment, combined with small but significant positive coefficients on cumulative UI registration and benefits, indicating a marginal increase in UI receipt relative to the reference.

The 3-4 p.m. slot displays an anomalous and internally inconsistent pattern: positive and significant coefficients on both employment outcomes, comparable in magnitude to the 7-8 a.m. result, alongside a null on cumulative UI registration and higher cumulative UI benefits. The explanation is selection into this slot, and it is visible on observables before any outcome is examined. Relative to the 9-10 a.m. reference, job-seekers seen at 3-4 p.m. are disproportionately already partially attached to the labour market: they are more often in interim employment in the meeting month (25% vs. 19%), more often enrolled in an active labour market programme (26% vs. 18%), and correspondingly more often off UI benefits (8.8% vs. 5.6%), while being indistinguishable on earnings, education, age, and tenure (Appendix Tables A10–A11). This is selection on current labour-market status rather than on human capital, and the same profile in far more extreme form marks the rare 5-6 p.m. slot; Section 2.7.1 further shows the anomaly is concentrated among single-spell individuals, consistent with first-time unemployed workers still cycling through interim work being scheduled into late slots. Such a group mechanically produces the pattern we observe here, positive on employment yet wrong-signed on UI outcomes: job-seekers already in interim work and off benefits behave differently on the UI margins regardless of meeting timing. It is this prior imbalance, confirmed by the internal incoherence of the outcomes rather than by employment coherence alone, that leads us to treat the 3-4 p.m. slot, like the rare 5-6 p.m. slot, as uninformative for identification purposes.

2.5. MAIN RESULTS: AVERAGE EFFECTS OF MEETING TIMING

93

Table 2.2: Effect of each meeting time slot on all outcomes (full sample)

	Employed at month 6			Cum. months employed		
	(1)	(2)	(3)	(1)	(2)	(3)
7–8 a.m.	0.00379*** (0.000432)	0.00234*** (0.000370)	0.00224*** (0.000369)	0.0275*** (0.00191)	0.0195*** (0.00153)	0.0191*** (0.00153)
11–12 a.m.	0.000534 (0.000712)	–0.0000336 (0.000606)	–0.000119 (0.000605)	0.0254*** (0.00315)	0.0168*** (0.00251)	0.0163*** (0.00250)
1–2 p.m.	–0.000251 (0.000368)	0.000108 (0.000315)	0.0000675 (0.000315)	0.000551 (0.00162)	0.00228* (0.00130)	0.00204 (0.00130)
3–4 p.m.	0.00348*** (0.000481)	0.00287*** (0.000412)	0.00279*** (0.000411)	0.0348*** (0.00214)	0.0233*** (0.00171)	0.0231*** (0.00171)
	Cum. months on UI			Cum. benefits paid		
	(1)	(2)	(3)	(1)	(2)	(3)
7–8 a.m.	–0.0200*** (0.00189)	–0.0107*** (0.00154)	–0.0102*** (0.00153)	–90.48*** (6.629)	–40.69*** (5.479)	–38.83*** (5.451)
11–12 a.m.	–0.0333*** (0.00314)	–0.0203*** (0.00254)	–0.0192*** (0.00252)	–109.1*** (11.17)	–57.33*** (9.153)	–54.38*** (9.098)
1–2 p.m.	0.00279* (0.00161)	0.00243* (0.00132)	0.00254* (0.00131)	14.86*** (5.571)	28.54*** (4.641)	28.78*** (4.618)
3–4 p.m.	–0.0126*** (0.00211)	–0.000808 (0.00172)	–0.000705 (0.00172)	–58.78*** (7.243)	17.87*** (5.991)	17.99*** (5.957)
	Income at month 6			Cum. months in ALMP		
	(1)	(2)	(3)	(1)	(2)	(3)
7–8 a.m.	14.88*** (2.337)	1.283 (2.016)	0.726 (2.009)	–0.00531*** (0.00149)	–0.0197*** (0.00138)	–0.0198*** (0.00137)
11–12 a.m.	1.347 (3.900)	–6.808** (3.308)	–7.246** (3.294)	0.000611 (0.00251)	–0.00883*** (0.00230)	–0.0103*** (0.00228)
1–2 p.m.	–4.965** (1.970)	–7.517*** (1.703)	–7.659*** (1.704)	–0.00397*** (0.00128)	–0.0106*** (0.00119)	–0.0103*** (0.00118)
3–4 p.m.	7.199*** (2.487)	–12.44*** (2.137)	–13.17*** (2.132)	–0.00123 (0.00169)	–0.0253*** (0.00156)	–0.0264*** (0.00155)
Observations	9,025,553	8,832,592	8,832,292	9,025,553	8,832,592	8,832,292

Notes: Standard errors clustered at the job-seeker level in parentheses. Reference category: 9–10 a.m. Columns (1), (2), (3) correspond to Individual FE, Individual×Spell FE, and Individual×Spell + Caseworker FE respectively. All outcomes are measured over the six months following the meeting; cumulative outcomes sum months or amounts over that window. “Employed at month 6” is a 0/1 indicator; benefits paid and income are in CHF; all other outcomes are in months. Full sample. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

2.6 Heterogeneity analysis

An average effect of meeting timing, however robust, says nothing about who benefits from it. If morning meetings work through channels that interact with job-seekers' constraints and circumstances (care obligations, distance from the labour market, the attention caseworkers devote to harder cases), the effect should differ across groups, and the differences are where the policy content lies. We therefore probe the average along several dimensions: gender, nationality, and proficiency in the language of the canton. In each case we estimate Versions 2 and 3 augmented with a full interaction between the time-slot dummies and the group indicator;⁷ because these characteristics are time-invariant, their main effects are absorbed by the individual $\times$ spell fixed effects, while the interactions remain identified from within-spell variation in meeting timing. One dimension stands out from this exercise, and it is the central finding of the chapter: the early-morning advantage is not distributed evenly by gender.

2.6.1 The gender interaction and its specificity

Table 2.3 reports the estimates for the gender interaction. In the preferred specification (V3), the 7–8 a.m. effect for men is +0.0015 and is significant; the interaction with the female indicator adds +0.0017, implying an effect for women of +0.0032, roughly twice as large. About half of the early-morning advantage experienced by women is thus gender-specific rather than common to all job-seekers, and the estimates are similar in V2 (column 1), indicating that caseworker-level heterogeneity plays no role in the pattern. The pattern is mirrored on the cumulative outcome (columns 3–4): the gap does not depend on which of the two employment outcomes one looks at.

As a check on inference rather than on the point estimate, we re-cluster the standard errors. Our default clusters at the job-seeker level; but because the treatment, the meeting hour, is assigned at the caseworker level, residuals may be correlated across meetings handled by the same caseworker even net of the caseworker fixed effect, which clustering by job-seeker would not capture. Re-clustering at the caseworker level, a more conservative choice given the far smaller number of clusters (roughly six thousand caseworkers against over a million job-seekers), leaves the gender interaction significant ($p = 0.032$).

In the same metric used for the average effect, the gender gap is substantial. Women in the sample have worse baseline outcomes than men (mean six-month employment of 0.367 vs. 0.396; mean cumulative employment of 1.46 vs. 1.65

⁷The rare 5–6 p.m. slot (~0.4% of meetings), flagged for selection in the descriptive analysis, is excluded from the heterogeneity analysis throughout.

2.6. HETEROGENEITY ANALYSIS

95

months), and it is for them that the morning meeting does more. For women, the 7–8 a.m. meeting adds +0.025 cumulative months of employment over the half-year, about three quarters of an extra day of work and 1.7% of the female mean, against +0.014 months (0.4 days, 0.9% of the male mean) for men. On point-in-time employment, the +0.32 percentage point effect for women is 0.9% of the female mean, roughly twice the corresponding relative figure for men (0.4%). The magnitudes remain modest in absolute terms; the point is their distribution, since the group with the weaker baseline outcomes, and the one least exposed to the early slot, is also the one that benefits most from it.

The most important feature of Table 2.3 is the specificity of the interaction across slots: it is significant only for 7–8 a.m., while the interactions for the 11–12 a.m., 1–2 p.m., and 3–4 p.m. slots on being employed 6 months after the meeting are all precise zeros. The one exception, which we report and do not interpret, is on the cumulative outcome: the 3–4 p.m. slot, flagged for selection in Section 2.5, carries a negative female interaction there (-0.0095 , $p = 0.005$), which we read under the same reservation as its base coefficients. Women benefit specifically from the early morning, not from non-standard meeting times in general. Section 2.6.3 discusses what this pattern implies for selection-based readings of the interaction.

2.6.2 Other dimensions of vulnerability

Gender is not the only dimension along which caseworkers might treat, or morning meetings might benefit, job-seekers differentially. A natural alternative is fragility in the labour market: caseworkers might devote more attention, or reserve particular slots, for job-seekers they perceive as harder to place, such as foreign nationals (43.5% of the analysis sample) or job-seekers with weak skills in the language of their canton. We re-estimate the interacted specification replacing the female indicator with, in turn, an indicator for foreign nationality and the categorical measure of local-language proficiency (Appendix Table A16). Neither dimension shows any heterogeneity on the focal slot: the 7–8 a.m. interaction with foreign nationality is a precise zero on both outcomes (-0.0009 on employment at six months, $+0.0001$ on cumulative months employed), and the interactions with limited and no language proficiency are likewise null on both outcomes, in both V2 and V3. A few significant interactions do appear for the intermediate proficiency group in non-focal slots (11–12 a.m. and 3–4 p.m., the latter the slot flagged for selection in Section 2.5); they do not touch the focal contrast and we do not interpret them. The early-morning effect is thus homogeneous across nationality and language groups, which sharpens the gender finding: among the dimensions we examine, gender is the only one along which the morning advantage varies. We

Table 2.3: Gender heterogeneity in the effect of meeting timing

	Employed at 6 months		Cumul. months employed	
	(1) V2	(2) V3	(3) V2	(4) V3
7–8 a.m. (base, men)	0.00159*** (0.000504)	0.00146*** (0.000503)	0.0150*** (0.00206)	0.0144*** (0.00205)
7–8 a.m. × female	0.00165** (0.000739)	0.00171** (0.000737)	0.0101*** (0.00307)	0.0105*** (0.00307)
11–12 a.m. × female	−0.000337 (0.00121)	−0.000282 (0.00121)	−0.00656 (0.00502)	−0.00604 (0.00500)
1–2 p.m. × female	−0.0000411 (0.000630)	−0.0000073 (0.000629)	0.000539 (0.00261)	0.000642 (0.00260)
3–4 p.m. × female	0.000292 (0.000817)	0.000430 (0.000816)	−0.0104*** (0.00342)	−0.00954*** (0.00341)
Implied effect for women, 7–8 a.m.	0.0032*** (0.0005)	0.0032*** (0.0005)	0.0251*** (0.0023)	0.0249*** (0.0023)
Individual × Spell FE	✓	✓	✓	✓
Caseworker FE		✓		✓
Observations	8,832,592	8,832,292	8,832,592	8,832,292

Notes: Each column is a single regression including the full set of time-slot dummies interacted with a female indicator, day-of-week, duration-category and meeting-type controls, caseworker workload controls, and a dummy for meeting the assigned caseworker. Reference slot: 9–10 a.m. The main effect of female is absorbed by the individual × spell fixed effects. All specifications include month-of-unemployment and the additional fixed effects of the main analysis; full sample. Standard errors clustered at the job-seeker level.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

2.6. HETEROGENEITY ANALYSIS

97

do not claim the list is exhaustive; other margins of heterogeneity (age, education, occupation) remain open for future work.

Because we probe several dimensions (gender, nationality, language) across several slots, the significant gender interaction at 7–8 a.m. raises a multiple-testing concern: with enough tests, one isolated rejection is expected by chance. We do not rely on a formal correction, because the pattern is not that of an isolated rejection: the effect is significant only at the focal slot and a precise zero at every other hour and along both other dimensions, it holds on two outcomes, doubles for women, and survives the within-spell switcher restriction (Section 2.6.3). A false positive generated by multiple testing would not reproduce this joint structure.

2.6.3 Selection into the morning slot by gender

The natural alternative reading of the interaction is selection: women observed at 7–8 a.m. may be a particular subgroup, better organised or less constrained, and their better outcomes would then reflect who they are rather than what the meeting time does. The following evidence speaks against this reading, and one candidate mechanism cannot be tested with our data; we take them in order. Overall women observed at 7–8 a.m. are observably indistinguishable from women observed at 9–10 a.m. Table 2.4 compares them across eleven characteristics. A joint test formally rejects balance ($F(11) = 99.1$), but with 2.1 million observations statistical rejection is uninformative about magnitude, and the differences are economically negligible: 0.13 years of age (37.89 vs. 37.76), 0.6 percentage points in child benefit receipt (8.4% vs. 9.1%), 0.8% in previous earnings (3,899 vs. 3,867 CHF), with the covariates jointly not explaining the slot assignment ($R^2 = 0.0005$). The same regression among men rejects as well ($F(11) = 46.7$, $R^2 = 0.0002$): observable sorting into the morning slot is mildly differential by gender but minuscule in both groups. There is no evidence of economically meaningful imbalance (which is not the same as no selection).

Additionally, the result survives the sharpest within-person checkfocusing on the so called *switchers*. Restricting the sample to spells featuring both a 7–8 a.m. meeting and a meeting in another slot (62.85% of individual-spells $N = 5,747,650$, so not a niche), the comparison is genuinely within-person: the same woman observed at different hours, across 844,453 individual-spell fixed effects and 665,701 job-seeker clusters. Table 2.5 reports the estimates: the interaction is unchanged, the effect for women is present as in the the full sample, and for men as well, and every other slot×female interaction is insignificant. The gender result therefore does not depend on comparisons between women who do and do not attend morning meetings. This is as a robustness to within-spell identification, not proof

of exogeneity: the switcher subsample is itself endogenous to meeting frequency.

The mechanism we cannot test directly is the role of children, which we initially tried to investigate. The morning-specific selection story, and the most plausible mechanism for the gender gap itself, both run through care obligations, so the natural next step would be to split the female effect by the presence of children. The register, however, contains no direct measure of children in the household; the closest variable is the receipt of child allowances, and it is a poor proxy for maternal care burdens, because allowances are typically paid to the household's main earner and thus captured disproportionately by men: only 8–9% of women in our sample receive them (Table 2.4), far below any plausible share of mothers. An exploratory specification interacting gender with this proxy showed no significant interaction with meeting timing, but given the measurement problem we read that as uninformative rather than as evidence against the care channel. Distinguishing the care mechanism from other sources of the gender gap therefore requires data we do not have, such as linked household registers, and we flag it as the main open question for future work.

Table 2.4: Balance on observables: women at 7–8 a.m. vs. 9–10 a.m.

	Women, 7–8 a.m.	Women, 9–10 a.m.
Age (years)	37.89	37.76
Tenure, last job (months)	20.92	21.01
Earnings, last job (CHF)	3,898.9	3,867.1
Ever self-employed	0.0379	0.0378
Ever on disability insurance	0.0179	0.0170
Ever on unemployment insurance	0.4306	0.4344
Child benefit receipt	0.0845	0.0908
Disability insurance receipt/application	0.0324	0.0349
Language skills (category)	1.140	1.137
Nationality (category)	2.171	2.209
Max. UI entitlement (days)	390.1	391.7
Observations	768,203	1,365,685
Joint test, women: $F(11) = 99.1$ ($p < 0.001$); $R^2 = 0.0005$		
Joint test, men: $F(11) = 46.7$ ($p < 0.001$); $R^2 = 0.0002$		

Notes: Means by time slot among women. The joint tests are from linear regressions of an indicator for the 7–8 a.m. slot (vs. 9–10 a.m.) on the eleven covariates, estimated separately by gender with robust standard errors ($N = 2,133,888$ women; $N = 2,426,396$ men). With samples of this size the tests reject mechanically; the economically relevant quantities are the magnitudes of the differences and the R^2 .

2.6. HETEROGENEITY ANALYSIS

99

Table 2.5: Gender heterogeneity on the switcher subsample (employment at six months)

	Switcher spells
7–8 a.m. (base, men)	0.00148*** (0.000535)
7–8 a.m. × female	0.00174** (0.000783)
11–12 a.m. × female	–0.000496 (0.00156)
1–2 p.m. × female	0.000281 (0.000812)
3–4 p.m. × female	0.000569 (0.00106)
Implied effect for women, 7–8 a.m.	0.00323*** (0.000574)
Individual × Spell FE	✓
Caseworker FE	✓
Observations	5,747,650
Clusters (job-seekers)	665,701

Notes: Sample restricted to individual-spells featuring both a 7–8 a.m. meeting and a meeting in another slot (62.85% of the full sample; 844,453 individual-spell fixed effects). Specification as in the preferred gender regression: full slot × female interactions, day-of-week, duration-category and meeting-type controls, caseworker workload controls, individual × spell, caseworker, month-of-unemployment and baseline cell fixed effects. Reference slot: 9–10 a.m. The 5–6 p.m. slot is excluded. The implied effect for women is the sum of the base and interaction coefficients, with the standard error computed by linear combination of the estimated coefficients. Standard errors clustered at the job-seeker level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

2.7 Robustness

2.7.1 The morning effect across single- and multi-spell individuals

For single-spell individuals, the individual fixed effect of Version 1 and the individual-by-spell fixed effect of Version 2 coincide, so the distinction between the two specifications is driven entirely by multi-spell individuals. On the spell-restricted sample, however, the effect of an early-morning meeting on point-in-time employment at month 6 falls to roughly one half of its full-sample magnitude and loses precision, while every cumulative outcome is basically unchanged. This subsection shows that the attenuation reflects sample composition and the mechanical properties of a point-in-time outcome, rather than a fragile effect.

We split the full sample according to whether an individual is observed with a single spell or with at least two spells, and we re-estimate the individual-by-spell (V2) specification separately on each group. In both groups identification remains within-person: the coefficient is estimated from individuals whose meetings within a spell occur at different times of day, and observations with no such variation are absorbed by the fixed effect⁸. Table 2.6 reports the coefficient on the 7–8 a.m. slot for all seven outcomes.

Two features of Table 2.6 carry the argument. The effect on cumulative employment is nearly identical across the two groups: an early meeting raises months employed over the half-year window by 0.021 for single-spell individuals and by 0.018 for multi-spell individuals, so the morning advantage is not weaker for individuals with a history of repeated unemployment. The effect on point-in-time employment, by contrast, diverges sharply: it is more than three times larger for single-spell individuals (0.0039 versus 0.0011).

This contrast is what one would expect when the treatment accelerates exit and the two groups differ in their attachment to employment. Point-in-time employment is a snapshot taken at a single date, whereas the cumulative outcome adds up time spent employed over the whole window. The pattern is consistent with a setting in which multi-spell individuals experience more frequent employment-to-unemployment transitions within the window: the morning advantage brings forward their exit from unemployment, which the cumulative measure records, but their subsequent transitions cause the point-in-time difference at month 6 to attenuate over time toward the control group. This reading is supported by the pattern of estimates, though the table does not itself measure the frequency of

⁸The split is an exact partition of the full sample: the two groups contain 3,952,426 and 4,880,166 observations respectively, and the multi-spell group reproduces the spell-restricted sample used elsewhere in the paper.

2.7. ROBUSTNESS

101

Table 2.6: Effect of a 7–8 a.m. meeting, by number of unemployment spells

Outcome	Single-spell (1 spell)	Multi-spell (≥ 2 spells)
Employed at month 6	0.0039*** (0.0005)	0.0011** (0.0005)
Cumulative months employed	0.0212*** (0.0023)	0.0182*** (0.0021)
Cumulative months registered	−0.0120*** (0.0023)	−0.0097*** (0.0020)
Cumulative UI benefits, CHF	−45.95*** (8.35)	−36.05*** (7.22)
Earnings at month 6, CHF	4.70 (3.11)	−1.36 (2.64)
Cumulative months in ALMP	−0.0172*** (0.0022)	−0.0218*** (0.0018)
Cumulative months interim earnings	−0.0086*** (0.0020)	−0.0113*** (0.0017)
Observations	3,952,426	4,880,166

Notes: Each cell reports the coefficient on the 7–8 a.m. slot (reference category: 9–10 a.m.) from a separate regression of the row outcome on the meeting-timing indicators, day-of-week, duration-category and meeting-type controls, and a caseworker-load control. All regressions include individual-by-spell, month-of-unemployment, and the baseline cell fixed effects (economic sector, occupation group, canton, calendar month of the meeting). Standard errors clustered by individual in parentheses. Single-spell: individuals observed with exactly one unemployment spell. Multi-spell: individuals observed with at least two spells. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

within-window transitions.

The two subsample coefficients account for the full-sample estimate⁹.

The split is informative about the anomalous afternoon slots. The selection patterns that motivate excluding the late slots are concentrated among single-spell individuals: for the latest slot, the single-spell coefficients on point-in-time employment and cumulative UI benefits are large and of opposite sign to a genuine treatment response, whereas the multi-spell coefficients are close to zero. This is consistent with selection into late scheduling being driven by first-time unemployed workers (for example those still serving notice or working part-time), providing additional support for our interpretation of the afternoon-slot anomalies as reflecting selection rather than treatment effects.

2.7.2 Placebo tests on non-realised meetings

For a meeting that did not take place at its scheduled time, the originally scheduled hour should not predict outcomes if what matters is the actual meeting time. We estimate the same V2/V3 specifications as the main analysis separately on postponed, not-attended, and cancelled meetings, using the originally scheduled hour as the placebo treatment. Two diagnostics support the design: the hourly distribution of postponed meetings matches that of realised meetings, so the scheduled hour is a genuine scheduling hour; and 574,607 postponed observations fall in individual-spells with at least two distinct meeting hours (192,509 clusters), sufficient for the within-spell specification. The estimation sample in Table 2.7 is somewhat smaller (502,772), because absorbing the individual $\times$ spell, month-of-unemployment and baseline cell fixed effects drops observations that become singletons once all these dimensions are jointly conditioned on.

Table 2.7 reports the postponed placebo for all seven outcomes. The pattern is identical across the two specifications; if anything, the placebo coefficient on point-in-time employment shrinks further once caseworker fixed effects are added (from 0.0016 to 0.0008). On point-in-time employment this estimate is small but imprecise (V2 standard error 0.0018), so its confidence interval contains both zero and our realised effect of 0.0022: the placebo is consistent with a null there but is not powered to exclude an effect of that magnitude, which is why the decisive test

⁹The N -weighted average of the two point-in-time coefficients, $(3,952,426 \times 0.0039 + 4,880,166 \times 0.0011) / 8,832,592 = 0.0024$, is close to the full-sample estimate of 0.0023; the two need not coincide exactly, since the pooled regression weights the group-specific effects by their share of the treatment variance rather than by sample size, and the near-equality indicates that these weights are similar across the two groups here. The smaller coefficient in the spell-restricted sample is thus essentially a composition effect, driven by the greater weight that sample places on multi-spell individuals, rather than a consequence of estimation. The cumulative outcomes, which do not display this divergence across groups, are stable across samples and remain our preferred summary of the effect.

2.7. ROBUSTNESS

103

Table 2.7: Postponed placebo: 7–8 a.m. coefficient, all outcomes

Outcome	Placebo (postponed)		Took place (V2)
	V2	V3	
Employed at month 6	0.0016 (0.0018)	0.0008 (0.0018)	0.0023*** (0.0004)
Cumulative months on UI	−0.0052 (0.0073)	−0.0015 (0.0073)	−0.0107*** (0.0015)
Cumulative UI benefits, CHF	5.76 (25.30)	15.85 (25.20)	−40.69*** (5.48)
Earnings at month 6, CHF	−7.00 (8.53)	−9.61 (8.53)	1.28 (2.02)
Cumulative months employed	0.0204*** (0.0073)	0.0174** (0.0073)	0.0195*** (0.0015)
Cumulative months in ALMP	−0.0186*** (0.0058)	−0.0179*** (0.0058)	−0.0197*** (0.0014)
Cumulative months interim earnings	−0.0256*** (0.0064)	−0.0234*** (0.0064)	−0.0101*** (0.0013)
Observations	502,772	502,438	8,832,592

Notes: Placebo treatment: originally scheduled hour (7–8 a.m. dummy, reference 9–10 a.m.) of postponed meetings, full postponed sample. V2: individual×spell, month-of-unemployment and baseline cell fixed effects; V3 adds caseworker fixed effects. Took-place column: the corresponding V2 estimate on realised meetings (Table 2.1). Standard errors clustered at the job-seeker level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

is the decomposition by source below rather than the bare null. The placebo is insignificant on point-in-time employment, cumulative months on UI, cumulative UI benefits, and monthly income, and significant on the three participation-intensity outcomes: cumulative months employed, months in active labour market programmes, and months in interim employment. The distinction is therefore not between point-in-time and cumulative outcomes, since cumulative months on UI and cumulative benefits are among those that show no placebo effect; it runs between the employment and UI outcomes on one side and the participation-intensity outcomes on the other. This selective pattern is more informative than a uniform null, which could reflect low power, or a uniformly significant result, which would indicate pervasive selection: the test has the power to detect something, and what it detects is confined to specific outcomes. A supporting check shows the significant coefficients are not mechanical: the correlation between the originally scheduled hour and the hour at which the meeting is actually re-held is low (0.16), so they cannot be attributed to the actual hour inheriting the scheduled one. The natural reading is that the postponement decision itself is selective: which job-seekers postpone which meetings is not random, though our data do not identify the specific mechanism behind it.

The other two groups of non-realised meetings complete the picture, with a caveat about power. The not-attended placebo delivers a coefficient on point-in-time employment of 0.00005 on $N = 263,721$, with a standard error of about 0.0021: on a sample roughly thirty times smaller than the main one the standard error is about six times larger, so the confidence interval comfortably contains both zero and our realised effect of 0.0022, and the implied minimum detectable effect (at 80% power) is around 0.006, nearly three times the effect we would want to rule out. The point estimate is reassuringly close to zero, though the test is not powered to call it a precise null. The cancelled sample collapses to $N = 10,209$ after absorbing the fixed effects; all coefficients are insignificant but with large standard errors, so these nulls reflect lack of power rather than evidence of no effect, and are reported for completeness with the sample size made explicit.

The decomposition by source of postponement allows us to discriminate between two selection stories with very different implications for our design. If the residual significance reflected selection into *scheduling* (stronger or weaker job-seekers systematically assigned to particular hours), the scheduled hour would predict outcomes regardless of who later postpones the meeting: the placebo would be significant on both sides of the decomposition, and our identification would be compromised. If it instead reflects selectivity in the *postponement decision* (which job-seekers postpone which meetings is not random), it should be concentrated among postponements initiated by the job-seeker and vanish among

2.7. ROBUSTNESS

105

postponements forced by caseworker unavailability, which we proxy by caseworkers with at least two postponements on the same day (41.7% of postponed meetings). The proxy could misclassify some caseworker-initiated postponements as job-seeker-side; this misclassification attenuates the contrast between the two groups, so the test is conservative. Splitting the postponed sample by source reduces the observations available for the within-spell estimation: a spell now contributes only if it contains at least two same-source postponed meetings at distinct hours, so spells whose postponements are of mixed source drop out of both subsamples, leaving 360,354 observations in the by-source estimation (Table 2.8) against the 502,772 of the pooled placebo. For the same reason the caseworker-side share of the estimation sample (129,141 of 360,354, or 35.8%) falls below its 41.7% share of all postponed meetings, since the source split absorbs proportionally more caseworker-side singletons.

Table 2.8: Postponed placebo by source of postponement: 7–8 a.m. coefficient

Outcome	V2		V3	
	Caseworker-side	Job-seeker-side	Caseworker-side	Job-seeker-side
Employed at month 6	−0.00368 (0.00364)	0.000311 (0.00266)	−0.00572 (0.00367)	−0.00116 (0.00267)
Cumulative months employed	0.00425 (0.0152)	0.0238** (0.0110)	−0.00160 (0.0152)	0.0196* (0.0110)
Cumulative months on UI	0.0125 (0.0151)	−0.00565 (0.0111)	0.0154 (0.0151)	0.000652 (0.0110)
Cumulative UI benefits, CHF	61.13 (52.95)	10.09 (37.67)	59.44 (53.17)	26.18 (37.52)
Earnings at month 6, CHF	−43.01** (17.81)	−5.629 (12.66)	−49.22*** (17.94)	−10.12 (12.66)
Cumulative months in ALMP	−0.00442 (0.0122)	−0.0175** (0.00862)	−0.00836 (0.0122)	−0.0164* (0.00858)
Cumulative months interim earnings	−0.0193 (0.0135)	−0.0296*** (0.00973)	−0.0119 (0.0135)	−0.0274*** (0.00969)
Observations	129,141	231,213	128,878	230,891

Notes: Placebo treatment: originally scheduled hour (7–8 a.m. dummy, reference 9–10 a.m.) of postponed meetings, estimated separately by source of postponement. Caseworker-side: postponements by caseworkers with at least two postponements on the same day (proxy for caseworker unavailability). Job-seeker-side: all remaining postponements. Same specifications and fixed effects as the main analysis. Standard errors clustered at the job-seeker level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 2.8 reports the 7–8 a.m. placebo coefficient separately for the two groups. On the caseworker side, the employment, UI, and participation outcomes are insignificant in both V2 and V3, and the significant coefficients of the pooled placebo reappear only on the job-seeker side. This is what we would expect if the residual selection detected by the pooled placebo sits in the postponement decision rather than in the assignment of the meeting hour: the postponement decision plays no role in the main analysis, which uses realised meetings, whereas selection into the scheduled hour itself, the concern that would threaten identification, would show

up on the caseworker side as well, and it does not. On cumulative months employed in particular, the caseworker-side coefficient is insignificant, so the significance of the pooled placebo on this outcome need not qualify the main estimate. One caseworker-side coefficient is significant, monthly income; its bearing is limited, since the effect of realised meeting timing on income is itself close to zero and insignificant (Table 2.1), and it plausibly reflects selection into caseworker-side postponement on earnings-related characteristics. That the employment and UI nulls hold across both specifications suggests the caseworker-side pattern is not an artefact of a single specification.

2.8 Discussion

2.8.1 Mechanism and interpretation

We cannot say *why* early-morning meetings work, because the data do not separate the candidate channels: better caseworker concentration, more attentive counselling, more credible monitoring, or differences in meeting content. Almost every intermediate outcome we observe (cumulative months in programmes, cumulative sanctions, cumulative interim months, cumulative benefits) is a counter that accumulates over months spent unemployed. When a morning meeting speeds exit, all of these counters fall simply because there are fewer unemployed months left in which to accumulate them, not because the caseworker acts differently at 7 a.m. The declines in cumulative programme participation and interim employment in Table 2.1 are consistent with this duration channel and do not point to a separate programme-assignment or monitoring channel. The only outcome free of the problem is a point-in-time flow, such as being sanctioned in the month of the meeting, but sanctions are rare and the test is underpowered. The chapter is therefore reduced-form. Candidate channels exist in the literature (monitoring and sanction threats that speed exit, or the attention and fatigue cycles documented in screening settings (Hangartner et al., 2021; Pope, 2016)), but our estimates do not favour one over another, and separating them would require data on meeting content that the register does not contain.

One story we can partly set aside is caseworker alertness, that the morning caseworker is simply fresher. It predicts a morning advantage that does not depend on who sits across the desk, which is hard to reconcile with the effect doubling for women (Section 2.6). That the gender gap appears at 7–8 a.m. and at no other hour points instead to constraints on the job-seeker's side: an early meeting can come before the demands of a day given over to care, whereas caseworker attention has no reason to be gender-specific to one slot. The near-zero income

effect fits the same reading (Section 2.5): the extra employment is concentrated in low-income jobs, closer to a search-activation channel (monitoring and immediate search actions that speed placement into whatever work is available) than to better matching. We read this with the caution the noisy income measure warrants. It does not pin down a mechanism, but it narrows the set worth testing once data on meeting content become available.

Several readings of the modest average are consistent with the evidence. A single meeting is a small dose, and the literature finds that meeting effects build over the sequence (Van den Berg et al., 2012); our treatment is the timing of one interaction within it. The effect is also concentrated, twice the average for women, so the average understates it rather than reflecting a uniformly weak treatment. And the six-month horizon is distant from the meeting, while the exit-timing proxies place the effect in the months just after it.

2.8.2 Policy implications

The gender result invites a discrimination reading, but our design does not support one. Discrimination would mean that equally qualified people are treated differently because of their group, through a biased decision-maker. What we show is different: a uniform administrative practice, standard scheduling, produces systematically different outcomes across groups. That is an *equity* finding. The returns to the most effective slot fall disproportionately to women, so how appointment times are allocated has distributional consequences even where no caseworker treats anyone differently. Whether caseworkers assign morning slots differentially by group, conditional on characteristics, is not identifiable here, because the identifying timing variation comes from scheduling that is exogenous to the job-seeker. The contribution is distributive, and the policy question is one of allocation.

Because schedules are an administrative choice, the findings support a nearly cost-neutral reallocation: giving early-morning slots to the groups with the highest returns, here women, a margin that rests on the clean specificity pattern of Section 2.6. It is implementable without new resources, and we offer it as a direction the estimates point to, not a fully evaluated policy.

One tempting inference is wrong. Multiplying the average effect by the volume of meetings overstates the return, because slots are a fixed resource and rescheduling is zero-sum: moving one job-seeker into the 7–8 a.m. slot moves another out. Under a homogeneous effect, any reallocation nets to zero, so the policy value lies entirely in the heterogeneity. Each swap that gives an early slot to a woman rather than a man yields about 0.17 percentage points on six-month employment

(0.0032 against 0.0015) and 0.0105 cumulative months of employment; per thousand swaps, roughly ten additional months of employment, at no budgetary cost. The room for it exists: women hold 44.7% of 7–8 a.m. meetings against 48.1% at 9–10 a.m. (Appendix Table A10), and no slot currently sorts on observables (Appendix Tables A10–A13), so a reallocation would displace no existing targeting.

Two limits bound the exercise. The estimates come from marginal variation in timing, and need not survive a large-scale reallocation that changes the composition of every slot; such a reallocation would also carry general-equilibrium effects, since someone has to take the afternoon slots, that fall outside the design. And enlarging the *supply* of early slots, rather than re-sorting within the existing one, is not free, since caseworkers would have to work earlier. Reallocation can capture the heterogeneity, not the average.

2.8.3 Limitations

We restate the limits in one place. Identification is not experimental: selection on unobservables (idiosyncratic within-spell motivation perceived by caseworkers, or morning-specific selection of unconstrained women) cannot be logically excluded, although the fixed-effects structure, the balance magnitudes, the switcher subsample, and the placebo pattern make it an unlikely driver. The mechanism is not separable with these data. The three participation-intensity outcomes triggered the pooled postponed placebo; the decomposition by source shows that this reflects selectivity in the postponement decision rather than in scheduling, and the caseworker-side placebo is a clean null (Section 2.7.2), but we flag these outcomes as the ones where the placebo evidence required this additional step. The intensity margin (how much time the caseworker devotes to each job-seeker, arguably the channel through which differential treatment across groups would operate most directly) remains outside the reach of these data, because recorded durations are available only in coarse categories rather than as a continuous measure (Section 2.3); it is the natural complement to the timing margin studied here, and pursuing it would require registers with finely measured durations. How the effect varies over the unemployment spell is a natural dimension of heterogeneity that we do not report. Our outcomes are measured six months after the meeting, while the exit-timing proxies suggest the effect operates in the months immediately following it; short-horizon outcomes, starting with employment one month after the meeting, are the first extension planned on the updated register data. And the gender heterogeneity, while robust and specific, is suggestive about *why* women benefit more: care constraints are a hypothesis we do not identify.

2.9 Conclusion

This chapter asked whether an apparently trivial choice in how public employment services are run, the hour at which mandatory caseworker meetings are scheduled, affects the labour market outcomes of the unemployed, and whether those effects fall evenly across job-seekers. Using more than 11 million meetings from the Swiss unemployment insurance registers and within-individual, within-spell variation in meeting timing, we find that timing matters, but modestly: early-morning meetings are associated with a 0.22 percentage point higher probability of employment six months later. What matters more than the size of this average is its distribution. The early-morning advantage is about twice as large for women, is specific to that slot, and holds up across every robustness check we ran, including the restriction to within-spell switchers.

The wider point is that how a public service is run matters beyond what the policy formally contains. Scheduling carries no distributive intent, yet it does not spread the returns to its most effective slot evenly across job-seekers. Meeting timing is a lever that employment services already use every day, at no extra cost, and our results show that how they use it shapes who gains from counselling. Showing that this dimension exists, that its average return is small, and that its benefits are uneven, is a first step toward treating scheduling as a deliberate choice. Extending the analysis to monthly horizons with the newly assembled data, and modelling exit as a duration process, are the natural next steps.

References

- Arni, Patrick, Gerard J. van den Berg, and Rafael Lalive (2020). "Treatment Versus Regime Effects of Carrots and Sticks". In: *Journal of Business & Economic Statistics*, pp. 1–17. DOI: 10.1080/07350015.2020.1784744. eprint: <https://doi.org/10.1080/07350015.2020.1784744>. URL: <https://doi.org/10.1080/07350015.2020.1784744>.
- Arni, Patrick and Amelie Schiprowski (2019). "Job search requirements, effort provision and labor market outcomes". In: *Journal of Public Economics* 169.C, pp. 65–88. DOI: 10.1016/j.jpubeco.2018.09. URL: <https://ideas.repec.org/a/eee/pubeco/v169y2019icp65-88.html>.
- Arni, Patrick P (2015). "Opening the Blackbox: How Does Labor Market Policy Affect the Job Seekers' Behavior? A Field Experiment". In.
- Behncke, Stefanie, Markus Frölich, and Michael Lechner (Apr. 2010a). "A Caseworker Like Me: Does the Similarity between Unemployed and Caseworker

- 129 Increase Job Placements?" In: *Economic Journal* 120(549), pp. 1430–1459. URL: <https://ideas.repec.org/p/iza/izadps/dp3437.html>.
- 12 Behncke, Stefanie, Markus Frölich, and Michael Lechner (2010b). "Unemployed and their caseworkers: should they be friends or foes?" In: *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 173.1, pp. 67–92. DOI: <https://doi.org/10.1111/j.1467-985X.2009.00600.x>. eprint: <https://rss.onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-985X.2009.00600.x>.
- 82 Berg, Gerard van den and Bas van der Klaauw (2006). "Counseling and Monitoring of Unemployed Workers: Theory and Evidence from a Controlled Social Experiment." In: *International Economic Review* 47.3, pp. 895–936. URL: <https://EconPapers.repec.org/RePEc:ier:iecrev:v:47:y:2006:i:3:p:895-936>.
- 7 Böheim, René, Rainer Eppel, Helmut Mahringer, et al. (2022). *More Caseworkers Shorten Unemployment Durations and Save Costs. Results from a Field Experiment in an Austrian Public Employment Office*. Tech. rep. WIFO.
- 15 Cederlöf, Jonas, Martin Söderström, and Johan Vikström (2025). "The role of caseworkers: Job finding, job quality and determinants of value-added". In: *Journal of the European Economic Association*, jvaf055.
- 60 Correia, Sergio and Noah Constantine (2026). "REGHDFE: Stata module to perform linear or instrumental-variable regression absorbing any number of high-dimensional fixed effects". In.
- 42 Danziger, Shai, Jonathan Levav, and Liora Avnaim-Pesso (2011). "Extraneous factors in judicial decisions". In: *Proceedings of the National Academy of Sciences* 108.17, pp. 6889–6892.
- 13 Dromundo Mokrani, Sofia (2021). "Three essays on labor market policies: job search assistance and digital tools". PhD thesis. Paris 1.
- 33 Federal Statistical Office (2016). "Gross Domestic Product per Capita". In: www.bfs.admin.ch.
- 96 Gaggero, Alessio and Denni Tommasi (2023). "Time of day and high-stake cognitive assessments". In: *The Economic Journal* 133.652, pp. 1407–1429.
- 4 Gerfin, Michael and Michael Lechner (Oct. 2002). "A Microeconomic Evaluation of the Active Labour Market Policy in Switzerland". In: *The Economic Journal* 112.482, pp. 854–893. ISSN: 0013-0133. DOI: [10.1111/1468-0297.00072](https://doi.org/10.1111/1468-0297.00072). eprint: <https://academic.oup.com/ej/article-pdf/112/482/854/26548131/ej0854.pdf>. URL: <https://doi.org/10.1111/1468-0297.00072>.
- 73 Haggag, Kareem, Richard W Patterson, Nolan G Pope, and Aaron Feudo (2021). "Attribution bias in major decisions: Evidence from the United States Military Academy". In: *Journal of Public Economics* 200, p. 104445.
- 47 Hangartner, Dominik, Daniel Kopp, and Michael Siegenthaler (2021). "Monitoring hiring discrimination through online recruitment platforms". In: *Nature* 589.7843, pp. 572–576.
- 116
- 7
- 130
- 70
- 10

REFERENCES

111

- Heissel, Jennifer A and Samuel Norris (2018). "Rise and shine: The effect of school start times on academic performance from childhood through puberty". In: *Journal of Human Resources* 53.4, pp. 957–992.
- Hervé, Justine, Subha Mani, Jere R Behrman, Ramanan Laxminarayan, and Arindam Nandi (2024). *Food Coma is Real: The Effect of Digestive Fatigue on Adolescents' Cognitive Performance*. Tech. rep. IZA Discussion Papers.
- Homrighausen, Pia and Michael Oberfichtner (2025). "Do caseworker meetings prevent unemployment? Evidence from a field experiment". In: *European Economic Review*, p. 105215.
- Huber, Martin, Michael Lechner, and Giovanni Mellace (2017). "Why do tougher caseworkers increase employment? The role of program assignment as a causal mechanism". In: *Review of Economics and Statistics* 99.1, pp. 180–183.
- Maibom, Jonas, Michael Rosholm, and Michael Svarer (2017). "Experimental Evidence on the Effects of Early Meetings and Activation". In: *The Scandinavian Journal of Economics* 119.3, pp. 541–570. DOI: <https://doi.org/10.1111/sjoe.12180>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/sjoe.12180>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/sjoe.12180>.
- Manning, Alan (2009). "You can't always get what you want: The impact of the UK Jobseeker's Allowance". In: *Labour Economics* 16.3, pp. 239–250. ISSN: 0927-5371. DOI: <https://doi.org/10.1016/j.labeco.2008.09.005>. URL: <https://www.sciencedirect.com/science/article/pii/S092753710800105X>.
- Morlok, M., D. Liehti, R. Lalive, A. Osikominu, and J. Zweimüller (2014). "Evaluation der arbeitsmarktlichen Massnahmen: Wirkung auf Bewerbungsverhalten und -chancen". In: *SECO Publikationen, Arbeitsmarktpolitik* 41.
- Oster, Emily (2019). "Unobservable selection and coefficient stability: Theory and evidence". In: *Journal of Business & Economic Statistics* 37.2, pp. 187–204.
- Pope, Nolan G (2016). "How the time of day affects productivity: Evidence from school schedules". In: *Review of Economics and Statistics* 98.1, pp. 1–11.
- Schiprowski, Amelie (2020). "The Role of Caseworkers in Unemployment Insurance: Evidence from Unplanned Absences". In: *Journal of Labor Economics* 38.4, pp. 1189–1225. DOI: [10.1086/706092](https://doi.org/10.1086/706092). eprint: <https://doi.org/10.1086/706092>. URL: <https://doi.org/10.1086/706092>.
- Schiprowski, Amelie, Julia Schmidtke, Johannes Schmieder, and Simon Trenkle (2024). "The effects of unemployment insurance caseworkers on job search effort". In: *AEA Papers and Proceedings*. Vol. 114. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 567–571.
- Shroff, Ravi and Konstantinos Vamvourellis (2022). "Pretrial release judgments and decision fatigue". In: *Judgment and Decision Making* 17.6, pp. 1176–1207.

73 State Secretary for Economic Affairs (SECO) (2017). "Die Lage auf dem Arbeitsmarkt im Februar 2017". In: www.seco.admin.ch.

4 Van den Berg, Gerard J., Iris Kesternich, Gerrit Müller, and Bettina M Siflinger (2020). "Reciprocity and the Interaction between the Unemployed and the Case-worker." In: *CEPR Discussion Paper* DP15038.

38 Van den Berg, Gerard J., Lene Kjaersgaard, and Michael Rosholm (2012). "To Meet or Not to Meet (Your Case Worker) – that is the Question." In: *IZA Discussion Paper* 6476.

3 Appendix 2.A Institutional background

122 Switzerland is a federal country with 26 cantons and three major language regions (French, German, and Italian). It is relatively wealthy, with GDP per capita of roughly 78,000 CHF and a low unemployment rate of 3 to 4% (State Secretary for Economic Affairs (SECO), 2017; Federal Statistical Office, 2016). Unemployment insurance is compulsory for workers. To obtain benefits, an unemployed person must register at the regional employment office (RAV) responsible for her community, be available for placement, search actively for suitable work, and take any steps the office deems necessary to improve her employability. Benefits are claimed at a branch of the unemployment insurance fund of the claimant's choice, which approves and pays them. Eligibility requires Swiss residency and the relevant work permits, together with at least twelve months of employment in the two years preceding the claim. Benefits replace about 70% of the previous wage, or 80% for claimants with dependent children or a dependent partner, up to a cap (see Behncke et al., 2010b); the potential duration depends on age, contribution history, and family situation, reaching a statutory maximum of 24 months. Services are free and open to anyone, not only to those entitled to benefits.

32 Registration is followed by the assignment of a caseworker, as in most OECD countries. According to Behncke et al. (2010b), the most common assignment criteria are caseload and occupation or industry sector, each named by roughly half of the caseworkers and officials polled, followed by randomisation. Regional offices enjoy considerable autonomy, in part a reflection of the federal structure. Caseworkers monitor job-search effort, can make binding agreements on search activity, refer job-seekers to vacancies, and decide on assignment to training and employment programmes on the basis of the job-seeker's employment history, subjective employability rating, and similar information; federal eligibility rules are deliberately broad, so that discretion rests largely with the office, and participation in a measure can effectively be compelled through the threat of sanctions. Beneficiaries are subject to compliance checks, with the control function exercised in part by

caseworkers and sanctions imposed by local offices or unemployment funds. On average a caseworker meets a job-seeker about once a month. Total yearly expenditure on Swiss active labour market policies exceeds 500 million CHF (Morlok et al., 2014), and Gerfin and Lechner (2002) classify these policies into training courses, employment programmes, and temporary employment schemes.

These meetings are governed by the Unemployment Insurance Act (AVIG) and its ordinance (AVIV). What matters for our design is that the office sets the interview dates for each insured person rather than leaving the choice to her (Art. 21(2) AVIV), and the case law is explicit on this point, holding that the job-seeker must organise herself so as to comply and may at most request a postponement (Administrative Court of the Canton of Bern, decision no. 200.2016.1229). Two channels nonetheless leave her some influence over timing: she may formally request that an assigned date be moved (Art. 25 AVIV), and a caseworker may informally agree the next slot at the end of a meeting, a practice that varies across offices. Office-set scheduling is therefore the institutional default rather than a guarantee of random assignment, which is why the identification of Section 2.4 does not rest on it alone: the placebo decomposition by source of postponement (Section 2.7.2) shows that the residual selection is concentrated in job-seeker-initiated postponements, the margin Art. 25 describes, and is absent on the caseworker side that supplies the identifying variation.

Appendix 2.B Baseline OLS analysis

This appendix documents the descriptive evidence from which the chapter started. The exercise differs from the analysis of Section 2.4.1 in two respects that we state at the outset. It uses no job-seeker fixed effects, so that the comparison is across job-seekers rather than within them; and its outcome is anchored to the spell rather than to the meeting, since Y_{im} below records employment six months after the start of the unemployment spell, whereas Y_{im}^6 in the main text records employment six months after meeting m . The two are different quantities, and we do not read the estimates below as effects. Indexing follows Section 2.4.1 throughout.

The model is a linear probability model estimated separately for each month of the spell:

$$Y_{im} = \beta_1 \text{time}_{im} + \beta_2 \text{day}_{im} + \beta_3 Z_{im} + \beta_4 X_{i,s(i,m)} + \beta_5 \text{share1}_{i,s(i,m)} + \beta_6 W_{c(i,m)} + \mu_{a(i,m)} + \gamma + \varepsilon_{im} \quad (2.4)$$

The unit of observation is a meeting, as in the main text, and the coefficient on the 7–8 a.m. slot is again the object of interest. The month of the spell $t(i, m)$ does not

appear in the equation because it selects the estimation sample rather than entering it: a separate regression is run for each of its values, so within each regression $t(i, m)$ is constant. This is the other role the same index plays in Section 2.4.1, where it varies across observations and is absorbed as a fixed effect.

Z_{im} collects the duration category and the type of the meeting, together with the indicator for whether it was conducted by the caseworker responsible for i rather than by a substitute. $X_{i,s(i,m)}$ collects job-seeker characteristics and employment history, here also including the job-seeker's status in the month of the meeting and the number of meetings already held; $\text{share1}_{i,s(i,m)}$ is the fraction of time spent in dependent employment before the current spell. $W_{c(i,m)}$ collects caseworker characteristics, $\mu_{a(i,m)}$ are fixed effects for the responsible caseworker, and γ collects fixed effects for economic sector, occupation group, canton, and calendar month of the year. Two further differences from the main analysis should be borne in mind when reading the tables: the sample is restricted to meetings with a recorded duration between 30 and 120 minutes, rather than the 20 to 120 window of Table A7; and standard errors are robust to heteroskedasticity but not clustered, so they do not account for the repeated appearance of the same job-seeker.

Results: first month of the spell. The day of the week shows no stable association with any outcome. The morning pattern that motivated the chapter, by contrast, is visible already here. Table A1 reports the coefficient on the 7–8 a.m. slot, adding controls sequentially. Without any controls it is small, negative and insignificant, which indicates that job-seekers seen in the early slot are negatively selected on observables. Once job-seeker characteristics enter, it turns positive and stays between 0.29 and 0.38 percentage points as employment history, caseworker characteristics and caseworker fixed effects are added; in column (5) it is 0.37 percentage points relative to 9–10 a.m. The stability across columns (2) to (5), and in particular the insensitivity to caseworker fixed effects, is what made the question worth pursuing. It is not evidence of an effect: the scheduling decision rests with the caseworker, and stability with respect to observables is not stability with respect to unobservables.

Table A2 shows the same coefficient for all outcomes under the specification of column (5). The pattern is internally consistent: early-morning meetings are associated with higher employment and labour income at both the six- and twelve-month horizons, and symmetrically with lower unemployment insurance receipt, benefit amounts, and registered unemployment.

2.B. BASELINE OLS ANALYSIS

115

Table A1: Coefficient on 7–8 a.m. dummy: sequential controls (first month, outcome: employment after 6 months)

	(1)	(2)	(3)	(4)	(5)
7–8 a.m.	–0.0011 (0.0011)	0.0038*** (0.0011)	0.0030*** (0.0010)	0.0029*** (0.0010)	0.0037*** (0.0011)
Job-seeker characteristics		✓	✓	✓	✓
Employment history			✓		✓
Caseworker characteristics				✓	✓
Caseworker fixed effects					✓
Substitute caseworker dummy					✓

Notes: Standard errors in parentheses. Column (5) is the preferred specification (equation 2.4).

Table A2: Effect of 7–8 a.m. meeting on all outcomes (first month, preferred specification)

Outcome	Coefficient
Employment after 6 months	0.0037*** (0.0011)
Employment after 12 months	0.0060*** (0.0011)
Income after 6 months (CHF)	19.45*** (6.316)
Income after 12 months (CHF)	24.45*** (6.347)
UI benefits after 6 months (CHF)	–13.40*** (3.674)
UI benefits after 12 months (CHF)	–16.81*** (3.475)
Receiving UI benefits after 6 months	–0.0038*** (0.0011)
Registered as unemployed after 6 months	–0.0035*** (0.0011)

Notes: Standard errors in parentheses. All outcomes relative to the 9–10 a.m. reference slot.

Results: months 2–5 of the spell. Tables A3–A6 repeat the exercise separately for the second to the fifth month of the spell, using the same sequential control structure. The morning coefficient is positive and significant in each of them, which is the point of the exercise: the regularity is not an artefact of looking only at the first month. We do not compare its magnitude across these tables. Because the outcome is fixed at six months from the start of the spell while the meeting moves later within it, the window separating meeting from outcome shrinks by one month from one table to the next, from six months in Table A1 to two in Ta-

ble A6; the ordering of the coefficients across tables therefore reflects the length of that window and cannot be read as the morning advantage growing with the stage of the spell. Unlike in the first month, the raw coefficients in column (1) are positive and larger than the controlled estimates, so that adding controls attenuates rather than reverses them; the estimates nonetheless remain stable once employment history is accounted for, and caseworker fixed effects again leave them nearly unchanged. Because this baseline cannot rule out time-varying unobservables driving both morning assignment and outcomes, all results in the main text rely on the fixed-effects specifications.

2.B. BASELINE OLS ANALYSIS

117

Table A3: Effect of 7–8 a.m. meeting: second month of unemployment

	(1)	(2)	(3)	(4)	(5)
Employment after 6 months	0.0118*** (0.0013)	0.0103*** (0.0013)	0.0077*** (0.0013)	0.0078*** (0.0013)	0.0072*** (0.0013)
Employment after 12 months	0.0141*** (0.0014)	0.0125*** (0.0013)	0.0098*** (0.0013)	0.0099*** (0.0013)	0.0084*** (0.0013)
Income after 6 months	84.84*** (8.097)	72.56*** (7.843)	44.83*** (7.361)	44.73*** (7.395)	43.44*** (7.458)
Income after 12 months	93.30*** (8.437)	80.86*** (8.088)	53.25*** (7.453)	53.34*** (7.500)	48.61*** (7.541)
UI benefits after 6 months	−31.88*** (4.845)	−29.93*** (4.652)	−26.70*** (4.286)	−27.49*** (4.305)	−24.57*** (4.340)
UI benefits after 12 months	−30.67*** (4.357)	−29.34*** (4.284)	−26.23*** (4.032)	−27.70*** (4.051)	−21.54*** (4.084)
Receiving UI after 6 months	−0.0113*** (0.0013)	−0.0097*** (0.0013)	−0.0071*** (0.0013)	−0.0074*** (0.0013)	−0.0060*** (0.0013)
Registered unemployed after 6 months	−0.0110*** (0.0013)	−0.0102*** (0.0013)	−0.0079*** (0.0013)	−0.0081*** (0.0013)	−0.0066*** (0.0013)
Job-seeker characteristics		✓	✓	✓	✓
Employment history			✓		✓
Caseworker characteristics				✓	✓
Caseworker fixed effects					✓
Substitute caseworker dummy					✓

Notes: Standard errors in parentheses.

CHAPTER 2. CASEWORKERS' MEETINGS WITH JOB-SEEKERS

Table A4: Effect of 7–8 a.m. meeting: third month of unemployment

	(1)	(2)	(3)	(4)	(5)
Employment after 6 months	0.0178** (0.0016)	0.0140*** (0.0016)	0.0093*** (0.0015)	0.0097*** (0.0015)	0.0080*** (0.0015)
Employment after 12 months	0.0139*** (0.0016)	0.0119*** (0.0016)	0.0073*** (0.0015)	0.0076*** (0.0015)	0.0053*** (0.0016)
Income after 6 months	125.9*** (9.657)	110.1*** (9.374)	54.10*** (8.806)	56.23*** (8.848)	49.17*** (8.955)
Income after 12 months	111.2*** (10.11)	105.4*** (9.706)	48.13*** (8.988)	49.32*** (9.031)	40.29*** (9.138)
UI benefits after 6 months	-43.88*** (5.884)	-28.01*** (5.640)	-35.46*** (5.167)	-36.29*** (5.193)	-29.68*** (5.266)
UI benefits after 12 months	-24.26*** (5.279)	-16.44*** (5.196)	-23.21*** (4.878)	-23.78*** (4.901)	-15.14*** (4.977)
Receiving UI after 6 months	-0.0171*** (0.0016)	-0.0126*** (0.0016)	-0.0093*** (0.0015)	-0.0096*** (0.0015)	-0.0070*** (0.0015)
Registered unemployed after 6 months	-0.0171*** (0.0016)	-0.0127*** (0.0016)	-0.0095*** (0.0015)	-0.0099*** (0.0015)	-0.0073*** (0.0015)
Job-seeker characteristics		✓	✓	✓	✓
Employment history			✓		✓
Caseworker characteristics				✓	✓
Caseworker fixed effects					✓
Substitute caseworker dummy					✓

Notes: Standard errors in parentheses.

2.B. BASELINE OLS ANALYSIS

119

Table A5: Effect of 7–8 a.m. meeting: fourth month of unemployment

	(1)	(2)	(3)	(4)	(5)
Employment after 6 months	0.0216*** (0.0017)	0.0177*** (0.0017)	0.0107*** (0.0016)	0.0112*** (0.0016)	0.0098*** (0.0017)
Employment after 12 months	0.0202*** (0.0018)	0.0180*** (0.0017)	0.0113*** (0.0017)	0.0116*** (0.0017)	0.0101*** (0.0017)
Income after 6 months	151.20*** (10.43)	135.60*** (10.14)	69.89*** (9.524)	73.38*** (9.570)	67.63*** (9.704)
Income after 12 months	134.50*** (11.43)	129.60*** (11.03)	63.70*** (10.23)	67.08*** (10.31)	61.09*** (10.57)
UI benefits after 6 months	−59.90*** (6.571)	−42.80*** (6.299)	−41.59*** (5.738)	−42.51*** (5.770)	−36.28*** (5.858)
UI benefits after 12 months	−37.05*** (5.900)	−27.67*** (5.806)	−29.22*** (5.434)	−29.87*** (5.464)	−24.95*** (5.552)
Receiving UI after 6 months	−0.0202*** (0.0018)	−0.0154*** (0.0017)	−0.0098*** (0.0017)	−0.0102*** (0.0017)	−0.0078*** (0.0017)
Registered unemployed after 6 months	−0.0201*** (0.0018)	−0.0152*** (0.0017)	−0.0097*** (0.0016)	−0.0101*** (0.0017)	−0.0073*** (0.0018)
Job-seeker characteristics		✓	✓	✓	✓
Employment history			✓		✓
Caseworker characteristics				✓	✓
Caseworker fixed effects					✓
Substitute caseworker dummy					✓

Notes: Standard errors in parentheses.

CHAPTER 2. CASEWORKERS' MEETINGS WITH JOB-SEEKERS

Table A6: Effect of 7–8 a.m. meeting: fifth month of unemployment

	(1)	(2)	(3)	(4)	(5)
Employment after 6 months	0.0189** (0.0018)	0.0157*** (0.0018)	0.0087*** (0.0018)	0.0090** (0.0018)	0.0071*** (0.0018)
Employment after 12 months	0.0195*** (0.0019)	0.0178*** (0.0019)	0.0108*** (0.0018)	0.0110*** (0.0018)	0.0082*** (0.0019)
Income after 6 months	137.40*** (10.95)	127.60*** (10.68)	59.88*** (10.05)	61.15*** (10.09)	50.65*** (10.27)
Income after 12 months	139.80*** (11.83)	139.30*** (11.42)	71.56*** (10.62)	72.19*** (10.67)	62.01*** (10.88)
UI benefits after 6 months	-55.99*** (7.177)	-37.28*** (6.887)	-36.84*** (6.223)	-38.11*** (6.259)	-28.93*** (6.362)
UI benefits after 12 months	-37.63*** (6.439)	-25.90*** (6.331)	-26.94*** (5.900)	-28.04*** (5.933)	-19.61*** (6.033)
Receiving UI after 6 months	-0.0190*** (0.0019)	-0.0147*** (0.0019)	-0.0089*** (0.0018)	-0.0093*** (0.0018)	-0.0061*** (0.0018)
Registered unemployed after 6 months	-0.0200*** (0.0019)	-0.0154*** (0.0019)	-0.0098*** (0.0018)	-0.0102*** (0.0018)	-0.0073*** (0.0018)
Job-seeker characteristics		✓	✓	✓	✓
Employment history			✓		✓
Caseworker characteristics				✓	✓
Caseworker fixed effects					✓
Substitute caseworker dummy					✓

Notes: Standard errors in parentheses.

Appendix 2.C Considered and discarded designs

Meetings recorded as postponed, cancelled, or not attended suggest alternative designs beyond their use as placebos. We considered three such designs and discarded all of them, for a common reason: *dynamic selection into non-realisation*. Meetings that are not realised are not a random sample. Job-seekers whose meetings are postponed or cancelled are positively selected on outcomes, since typically those already exiting unemployment are the ones whose meetings are dropped, so comparing outcomes between realised and non-realised meetings confounds the meeting effect with selection into skipping it.

The most promising variant is a rescheduling design, which uses postponed meetings that are subsequently held and treats the hour of the *re-scheduled* meeting as quasi-random, on the premise that the new hour is set by calendar availability rather than chosen by the job-seeker. Feasibility diagnostics supported this premise: the correlation between the scheduled and the re-held hour is 0.156, and on 380,927 re-scheduled meeting-slots a balance regression of “re-held in the morning” on observables yields $R^2 = 0.0007$, so observables explain 0.07% of the variation, with detectable differences that are economically negligible. The design nonetheless fails on two counts. The rescheduled sample is selected on *being postponed*, which is itself outcome-relevant; and the clean variant, which restricts to postponements plausibly initiated by the caseworker (proxied by caseworkers with multiple same-day postponements, 41.7% coverage), collapses the sample to roughly 52,000 observations, where standard errors triple and the effect is undetectable. The design is conceptually valid and empirically dead, right in theory but underpowered in practice, and we document it here so that it is not re-opened.

The remaining two designs, which compare realised meetings to cancelled (status 5) or not-attended (status 4) ones, fail on the same selection logic, compounded for cancelled meetings by the collapse of the within-spell sample ($N = 10,209$). The parallel data builds exist in the pipeline (06_dataprep.do, filtering by status) and are retained only for the placebo use of Section 2.7.2.

Appendix 2.D Intermediate outcomes and the duration contamination

We examined the available intermediate outcomes as candidate channels: cumulative months in active labour market programmes, cumulative sanctions (both the count of sanctions and the cumulative number of sanction days), cumulative months in interim work, cumulative months on UI, cumulative UI benefits,

and income. All cumulative counters share the structural problem described in Section 2.8: they accumulate over months spent unemployed, so any effect that accelerates exit reduces them mechanically. For example, cumulative months in programmes declines because a job-seeker who exits earlier spends fewer months exposed to programme assignment, not because caseworkers assign fewer measures in the morning. The same holds for cumulative sanctions.

The exception is the point-in-time flow variable recording whether a job-seeker was sanctioned in the month of the meeting: this is an act of the caseworker that precedes the outcome rather than a state that employment switches off, and it would constitute a genuine monitoring channel if it responded to meeting hour. Sanctions are however rare, and the test is underpowered; a null is uninformative. We report this as the boundary of what the data can say about mechanism.

Appendix 2.E Appendix tables

Table A7: Sample construction

	Meetings	Dropped
Full UI register, 2001–2018	20,425,634	
took place, at plausible hours (07:00–19:00)	16,083,719	4,341,915
plausible recorded duration (20–120 min): <i>analysis sample</i>	11,745,735	4,337,984
Estimation sample, individual FE	9,025,553	2,720,182
Estimation sample, individual $\times$ spell FE	8,832,592	192,961

Notes: Restrictions are applied cumulatively; “Dropped” is the reduction relative to the row above. The full register covers 1,259,358 job-seekers and 7,397 caseworkers. Meetings that did not take place (postponed, cancelled, not attended) are reserved for the placebo analysis of Section 2.7.2. The estimation-sample rows drop observations that become singletons once the fixed effects are absorbed; the individual $\times$ spell sample requires at least two meetings per spell at distinct hours.

Table A8: Description of control variables: job-seeker characteristics

Variable	Definition
<i>Status in the month of the meeting</i>	
Off unemployment insurance (ui0)	Not receiving UI benefits in the current month
Interim employment (zv0)	Holds an interim job in the current month
Programme participation (amm0)	Enrolled in an active labour market programme in the current month
Sanctioned (sanc0)	Sanctioned in the current month
<i>Sociodemographic characteristics and benefits</i>	
Woman (frau)	Indicator for female job-seekers
Age (alter)	Age in years; also entered grouped (alter_cat)
Ever self-employed (any2)	Any past self-employment spell
Ever on disability insurance (any3)	Any past DI spell
Ever on unemployment insurance (any4)	Any past UI spell
Other spells (any5)	Any past spell in other labour market states
Mobility (mobil)	High willingness to commute or relocate
Child allowance (kinderzulage)	Receives additional benefits for children
Disability insurance (iv)	Receives or has applied for DI
Training allowance (ausbildungszulage)	Receives additional benefits for training
Previous earnings (verdienst)	Earnings in the last job
Previous tenure (tenure)	Tenure in the last job
<i>Categorical variables</i>	
Education (ausbild)	Highest education level attained
Hierarchical position (code_funktion_n)	Hierarchical function in the last job, grouped
Language skills	Proficiency in the language of the canton, grouped
Nationality (nation)	Nationality, grouped
Entry status (type)	Labour market status at the beginning of the spell
Civil status (zivil)	Marital status
Last job workload (pensum)	Employment level of the last job, grouped
Benefit entitlement (anspruch)	Maximum UI benefit duration, in days
Meetings so far (anzahl)	Number of meetings held so far, grouped

Table A9: Description of control variables: employment history and caseworker characteristics

Variable	Definition
<i>Employment history</i>	
Pre-spell employment share (share1)	Share of time in dependent employment before the current spell
Average earnings, 60 months (din1_60)	Average monthly earnings over the past 60 months
Employment spells, 60 months (anz1_60)	Number of employment spells over the past 60 months
Employment duration, 60 months (des1_60)	Average duration of employment over the past 60 months
UI spells, 60 months (anz4_60)	Number of UI spells over the past 60 months
UI duration, 60 months (des4_60)	Average duration on UI over the past 60 months
Share out of the labour force (fes0_24)	Share of time out of the labour force over the past 24 months
Any OLF spell (any0_24)	Any spell out of the labour force in the past 24 months
Recent self-employment (any2_24)	Any self-employment in the past 24 months
Recent DI spell (any3_24)	Any DI spell in the past 24 months
Recent UI spell (any4_24)	Any UI spell in the past 24 months
Share on UI (fes4_24)	Share of time on UI over the past 24 months
<i>Caseworker and meeting characteristics</i>	
Assigned caseworker (dummy_caseworker)	Meeting held by the job-seeker's assigned caseworker
Current caseload (curcases)	Number of meetings conducted in the current month
Average meeting duration (dauer2)	Average recorded duration of the caseworker's meetings
Missing durations (mdauer2)	Share of the caseworker's meetings with missing recorded duration
Work experience (experience)	Caseworker's work experience since 2000
Total caseload (allcases)	Total number of meetings conducted since 2000
Average caseload (avcases)	Average number of meetings per month since 2000

2.E. APPENDIX TABLES

125

Table A10: Descriptive characteristics on control 1

Variables	(7–8)	(9–10)	(11–12)	(13–14)	(15–16)	(17–18)
Ui 0	0.0651 (0.2469)	0.0559 (0.2298)	0.0721 (0.2586)	0.0600 (0.2374)	0.0878 (0.2830)	0.2886 (0.4531)
Zv 0	0.2176 (0.4126)	0.1949 (0.3961)	0.2075 (0.4055)	0.2095 (0.4070)	0.2514 (0.4338)	0.4218 (0.4939)
Amm 0	0.2050 (0.4037)	0.1801 (0.3843)	0.2261 (0.4183)	0.2022 (0.4016)	0.2634 (0.4405)	0.2281 (0.4196)
Sanc 0	0.1238 (0.3294)	0.1118 (0.3151)	0.1073 (0.3095)	0.1133 (0.3169)	0.1063 (0.3083)	0.1591 (0.3657)
Frau	0.4469 (0.4972)	0.4811 (0.4996)	0.4752 (0.4994)	0.4882 (0.4999)	0.4625 (0.4986)	0.3307 (0.4705)
Alter	38.3678 (10.7897)	38.2177 (10.7071)	38.0212 (10.6909)	38.0638 (10.6778)	37.9590 (10.6674)	37.3669 (10.8056)
Any 2	0.0488 (0.2155)	0.0482 (0.2142)	0.0449 (0.2071)	0.0458 (0.2090)	0.0433 (0.2035)	0.0381 (0.1915)
Any 3	0.0258 (0.1585)	0.0254 (0.1575)	0.0232 (0.1505)	0.0256 (0.1578)	0.0249 (0.1558)	0.0267 (0.1612)
Any 4	0.4589 (0.4983)	0.4592 (0.4983)	0.4532 (0.4978)	0.4544 (0.4979)	0.4550 (0.4980)	0.4587 (0.4983)
Any 5	0.5773 (0.4940)	0.5833 (0.4930)	0.5907 (0.4917)	0.5908 (0.4917)	0.5882 (0.4922)	0.5803 (0.4935)
Mobil	0.0400 (0.1957)	0.0390 (0.1935)	0.0400 (0.1954)	0.0375 (0.1901)	0.0370 (0.1889)	0.0512 (0.2205)
Kinderzulage	0.1391 (0.3460)	0.1381 (0.3450)	0.1330 (0.3391)	0.1342 (0.3408)	0.1403 (0.3473)	0.1709 (0.3764)
Iv	0.0352 (0.1842)	0.0384 (0.1923)	0.0327 (0.1778)	0.0392 (0.1941)	0.0359 (0.1860)	0.0250 (0.1560)
Ausbildungszulage	0.0410 (0.1983)	0.0389 (0.1934)	0.0373 (0.1896)	0.0377 (0.1905)	0.0393 (0.1944)	0.0472 (0.2121)
Verdienst	4674.733 (2145.027)	4604.446 (2156.737)	4702.243 (2187.183)	4568.35 (2139.299)	4553.903 (2051.806)	4713.794 (1798.772)
Tenure	20.2343 (27.577)	20.34979 (27.518)	20.53654 (27.3457)	20.36244 (27.554)	20.47321 (27.594)	18.45606 (25.071)

Notes: Columns refer to meeting hour. Standard deviations in parentheses.

Table A11: Descriptive characteristics on control 2

Variables	(7–8)	(9–10)	(11–12)	(13–14)	(15–16)	(17–18)
Ausbild	2.3987 (1.5616)	2.4114 (1.5868)	2.5084 (1.6527)	2.3901 (1.5895)	2.3118 (1.5339)	2.1149 (1.3166)
Alter_cat	2.8386 (1.1425)	2.8237 (1.1338)	2.8049 (1.1312)	2.8077 (1.1306)	2.7970 (1.1294)	2.7333 (1.1486)
Code_funktion_n	22.2316 (0.5632)	22.2351 (0.5606)	22.2120 (0.5447)	22.2537 (0.5625)	22.2845 (0.5646)	22.2872 (0.5521)
Language	1.1488 (0.4518)	1.1448 (0.4445)	1.1378 (0.4317)	1.1577 (0.4575)	1.1799 (0.4855)	1.1929 (0.5163)
Nation	2.3084 (2.1425)	2.3239 (2.1540)	2.3698 (2.1920)	2.3541 (2.1807)	2.4426 (2.2218)	2.5273 (2.1946)
Type	1.3487 (0.7740)	1.3471 (0.7473)	1.3358 (0.7656)	1.3551 (0.7634)	1.3564 (0.8005)	1.3744 (0.9067)
Zivil	1.7284 (0.6623)	1.7304 (0.6615)	1.7188 (0.6678)	1.7304 (0.6607)	1.7349 (0.6539)	1.7008 (0.6425)
Pensum	3.4748 (0.9658)	3.4371 (0.9973)	3.4784 (0.9595)	3.4404 (0.9926)	3.4921 (0.9505)	3.6795 (0.7710)
Anspruch	390.259 (78.874)	391.137 (78.438)	390.647 (81.178)	389.978 (78.389)	391.070 (76.559)	396.166 (70.854)
Anzahl	4.8348 (3.7888)	4.8658 (3.6962)	5.3581 (3.6647)	4.9973 (3.7250)	5.6157 (3.7711)	6.1680 (3.9688)

Notes: Columns refer to meeting hour. Standard deviations in parentheses.

2.E. APPENDIX TABLES

127

Table A12: Descriptive characteristics on control 3

Variables	(7–8)	(9–10)	(11–12)	(13–14)	(15–16)	(17–18)
Din1_60	5249.018 (4787.854)	5154.778 (4548.963)	5238.328 (4660.920)	5093.329 (4450.432)	5030.303 (4270.310)	5022.517 (4165.745)
Anz1_60	2.1525 (1.4039)	2.1474 (1.4065)	2.1423 (1.4009)	2.1381 (1.4027)	2.1491 (1.4071)	2.2847 (1.5054)
Des1_60	24.610 (20.200)	24.681 (20.200)	24.643 (20.207)	24.542 (20.147)	24.447 (20.177)	23.513 (19.786)
Anz4_60	0.7953 (1.2426)	0.7937 (1.2448)	0.7774 (1.2272)	0.7795 (1.2263)	0.7901 (1.2402)	0.8762 (1.3590)
Des4_60	2.6828 (4.5986)	2.700 (4.6300)	2.7286 (4.6970)	2.6771 (4.6194)	2.6803 (4.6262)	2.3149 (4.0983)
Fes0_24	0.2333 (0.3037)	0.2334 (0.3038)	0.2314 (0.3035)	0.2369 (0.3052)	0.2354 (0.3041)	0.2358 (0.2972)
Any0_24	0.5707 (0.4950)	0.5712 (0.4950)	0.5679 (0.4954)	0.5773 (0.4940)	0.5742 (0.4945)	0.5901 (0.4918)
Any2_24	0.0266 (0.1610)	0.0261 (0.1594)	0.0240 (0.1530)	0.0245 (0.1547)	0.0226 (0.1487)	0.0190 (0.1367)
Any3_24	0.0167 (0.1280)	0.0166 (0.1276)	0.0153 (0.1228)	0.0168 (0.1285)	0.0163 (0.1266)	0.0160 (0.1254)
Any4_24	0.2565 (0.4367)	0.2537 (0.4351)	0.2480 (0.4318)	0.2488 (0.4323)	0.2487 (0.4322)	0.2777 (0.4479)
Fes4_24	0.0634 (0.1368)	0.0628 (0.1362)	0.0613 (0.1349)	0.0614 (0.1351)	0.0612 (0.1345)	0.0663 (0.1369)

Notes: Columns refer to meeting hour. Standard deviations in parentheses. Variable definitions in Table A9.

Table A13: Descriptive characteristics on control 4

Variables	(7–8)	(9–10)	(11–12)	(13–14)	(15–16)	(17–18)
Curcases	85.857 (39.893)	85.989 (38.975)	90.379 (36.663)	86.877 (38.092)	89.480 (37.516)	96.578 (43.215)
Dauer 2	75.448 (4870.827)	72.562 (4731.507)	60.002 (3006.860)	76.964 (5003.085)	76.876 (5047.187)	43.420 (472.591)
Mdauer2	0.0571 (0.1289)	0.0578 (0.1276)	0.0489 (0.1184)	0.0604 (0.1293)	0.0570 (0.1229)	0.0377 (0.0985)
Experience	78.020 (48.953)	78.523 (48.732)	78.310 (50.456)	77.308 (48.803)	75.436 (49.010)	65.560 (43.724)
All cases	4963.683 (4187.151)	5000.384 (4181.493)	5002.072 (4306.296)	4923.005 (4155.067)	4809.265 (4131.816)	4606.591 (4067.907)
Av. cases	62.576 (31.586)	62.737 (31.037)	62.622 (30.015)	62.682 (30.757)	62.901 (30.484)	68.412 (32.587)
Exp_cat	2.6579 (0.9804)	2.6700 (0.9728)	2.6491 (0.9995)	2.6444 (0.9820)	2.6005 (0.9878)	2.4287 (0.9246)

Notes: Columns refer to meeting hour. Standard deviations in parentheses. Average meeting duration (Dauer 2) is computed on the caseworker's raw recorded durations, before the duration cleaning applied to the analysis sample; its large dispersion reflects untrimmed extreme values and is not informative about actual meeting length.

2.E. APPENDIX TABLES

129

Table A14: Duration of implausible meetings.

Duration in minutes	Freq. (1)	Percent (2)	Cum. (3)
0–1	1,818,368	86.86	86.86
1–5	173,617	8.29	95.16
5–10	29,125	1.39	96.55
10–15	30,259	1.45	98.00
15–20	15,729	0.75	98.75
20–25	1,191	0.06	98.80
25–30	137	0.01	98.81
30–35	16,658	0.80	99.61
35–40	23	0.00	99.61
40–45	600	0.03	99.64
45–50	842	0.04	99.68
50–55	42	0.00	99.68
55–60	24	0.00	99.68
60–65	5,078	0.24	99.92
65–70	6	0.00	99.92
70–75	9	0.00	99.92
75–80	14	0.00	99.92
80–85	3	0.00	99.92
90–95	51	0.00	99.93
105–110	2	0.00	99.93
120–125	34	0.00	99.93
130–135	2	0.00	99.93
135–140	1	0.00	99.93
150–155	3	0.00	99.93
180–185	7	0.00	99.93
999 and over	1,506	0.07	100.00
Total	2,093,331	100.00	

Notes: The sample is composed of implausible meetings (between 7 p.m. and 7 a.m.). Cases with zero frequency not reported.

Table A15: Duration and status of implausible meetings

Duration (min)	Planned	Took place	Postponed	Not attended	Cancelled	Total
	(1)	(2)	(3)	(4)	(5)	(6)
0-1	3,206	1,764,647	6,302	44,213	0	1,818,368
1-5	1,475	161,059	6,839	3,240	1,004	173,617
5-10	110	28,021	106	870	18	29,125
10-15	136	29,433	135	545	10	30,259
15-20	108	14,935	198	426	62	15,729
20-25	8	1,137	30	15	1	1,191
25-30	0	117	1	19	0	137
30-35	53	15,506	415	624	60	16,658
35-40	0	23	0	0	0	23
40-45	0	587	3	4	6	600
45-50	1	685	102	41	13	842
50-55	0	42	0	0	0	42
55-60	0	24	0	0	0	24
60-65	17	4,857	63	128	13	5,078
65-70	0	6	0	0	0	6
70-75	0	9	0	0	0	9
75-80	0	13	1	0	0	14
80-85	0	3	0	0	0	3
90-95	1	46	2	2	0	51
105-110	0	2	0	0	0	2
120-125	1	31	1	0	1	34
130-135	0	1	0	1	0	2
135-140	0	1	0	0	0	1
150-155	0	3	0	0	0	3
180-185	0	7	0	0	0	7
999+	20	1,383	65	37	1	1,506
Total	5,136	2,022,578	14,263	50,165	1,189	2,093,331

Notes: Sample composed of implausible meetings (between 7 p.m. and 7 a.m.). Cases with zero frequency not reported.

2.E. APPENDIX TABLES

131

Table A16: Heterogeneity by nationality and language proficiency

	emp6		cemp6	
	(1) V2	(2) V3	(3) V2	(4) V3
<i>Panel A: Foreign nationality</i>				
7–8 a.m. (base, Swiss)	0.00331* (0.00183)	0.00308* (0.00183)	0.0196*** (0.00729)	0.0191*** (0.00727)
7–8 a.m. × foreign	−0.00101 (0.00187)	−0.000875 (0.00187)	−0.000124 (0.00746)	0.0000777 (0.00744)
11–12 a.m. × foreign	0.000911 (0.00332)	0.00106 (0.00332)	−0.0129 (0.0135)	−0.0120 (0.0135)
1–2 p.m. × foreign	−0.00178 (0.00161)	−0.00164 (0.00161)	0.00217 (0.00643)	0.00254 (0.00641)
3–4 p.m. × foreign	−0.00383* (0.00215)	−0.00374* (0.00215)	−0.0164* (0.00877)	−0.0160* (0.00875)
<i>Panel B: Language proficiency (base: at least good)</i>				
7–8 a.m. (base)	0.00240*** (0.000393)	0.00230*** (0.000392)	0.0192*** (0.00163)	0.0189*** (0.00162)
7–8 a.m. × limited	−0.000444 (0.00137)	−0.000450 (0.00137)	0.00189 (0.00583)	0.00140 (0.00582)
7–8 a.m. × none	−0.000742 (0.00194)	−0.000887 (0.00193)	0.00393 (0.00791)	0.00405 (0.00788)
11–12 a.m. × limited	−0.00721*** (0.00226)	−0.00714*** (0.00225)	−0.0506*** (0.00942)	−0.0499*** (0.00939)
11–12 a.m. × none	−0.00329 (0.00332)	−0.00316 (0.00332)	0.00512 (0.0135)	0.00577 (0.0134)
1–2 p.m. × limited	0.000372 (0.00115)	0.000431 (0.00115)	0.000201 (0.00484)	0.000747 (0.00483)
1–2 p.m. × none	0.000486 (0.00167)	0.000573 (0.00167)	0.00725 (0.00685)	0.00786 (0.00683)
3–4 p.m. × limited	−0.00979*** (0.00146)	−0.00978*** (0.00145)	−0.0575*** (0.00617)	−0.0574*** (0.00614)
3–4 p.m. × none	−0.00177 (0.00209)	−0.00185 (0.00209)	−0.00551 (0.00861)	−0.00557 (0.00859)
Individual × Spell FE	✓	✓	✓	✓
Caseworker FE		✓		✓
Observations	8,832,592	8,832,292	8,832,592	8,832,292

Notes: Each column of each panel is a single regression including the full set of time-slot dummies interacted with the group indicator, day-of-week, duration-category and meeting-type controls, caseworker workload controls, and a dummy for meeting the assigned caseworker; specification identical to Table 2.3. Foreign is defined as any nationality other than Swiss (43.5% of the sample). Language proficiency refers to skills in the language of the job-seeker's canton, in three categories (at least good; less than good; no knowledge). Main effects of the group indicators are absorbed by the individual × spell fixed effects. Reference slot: 9–10 a.m. The 5–6 p.m. slot is excluded. Standard errors clustered at the job-seeker level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A17: Status of plausible meetings by duration.

Duration (min)	Planned	Took place	Postponed	Not attended	Cancelled	Total
	(1)	(2)	(3)	(4)	(5)	(6)
0–1	2,253	568,728	6,085	39,839	0	616,905
	2.37%	3.54%	0.50%	4.95%	0.00%	3.37%
1–5	9,065	849,174	37,249	23,309	1,244	920,041
	9.53%	5.28%	3.07%	2.90%	0.93%	5.02%
5–10	3,614	660,222	6,049	14,605	329	684,819
	3.80%	4.10%	0.50%	1.81%	0.25%	3.74%
10–15	3,369	601,956	3,010	13,061	190	621,586
	3.54%	3.74%	0.25%	1.62%	0.14%	3.39%
15–20	4,864	1,143,827	13,724	26,430	1,823	1,190,668
	5.12%	7.11%	1.13%	3.28%	1.37%	6.50%
20–25	734	224,064	5,644	7,158	496	238,096
	0.77%	1.39%	0.46%	0.89%	0.37%	1.30%
25–30	221	102,778	2,053	3,166	124	108,342
	0.23%	0.64%	0.17%	0.39%	0.09%	0.59%
30–35	46,753	7,553,210	779,590	459,624	91,996	8,931,173
	49.17%	46.96%	64.16%	57.10%	69.10%	48.72%
35–40	85	33,369	697	303	160	34,614
	0.09%	0.21%	0.06%	0.04%	0.12%	0.19%
40–45	278	65,779	3,581	1,813	1,031	72,482
	0.29%	0.41%	0.29%	0.23%	0.77%	0.40%
45–50	4,331	1,083,691	97,259	62,702	14,587	1,262,570
	4.55%	6.74%	8.00%	7.79%	10.96%	6.89%
50–55	58	25,276	646	802	88	26,870
	0.06%	0.16%	0.05%	0.10%	0.07%	0.15%
55–60	88	20,632	599	580	72	21,971
	0.09%	0.13%	0.05%	0.07%	0.05%	0.12%
60–90	17,534	2,829,981	234,574	140,029	19,136	3,241,254
	18.44%	17.60%	19.31%	17.40%	14.37%	17.68%
90+	1,213	229,421	19,495	8,794	1,514	260,437
	1.28%	1.43%	1.60%	1.09%	1.14%	1.42%
Total	95,084	16,083,719	1,214,985	804,942	133,143	18,331,873

Notes: Sample restricted to plausible meetings (7 a.m.–7 p.m.). First row: counts; second row: percentages.

2.E. APPENDIX TABLES

133

Table A18: Status of plausible and implausible meetings.

Status	Implausible (1)	Plausible (2)	Total (3)
Planned	5,136 5.12%	95,084 94.88%	100,220 100.00%
Took place	2,022,578 11.17%	16,083,719 88.83%	18,106,297 100.00%
Postponed	14,263 1.16%	1,214,985 98.84%	1,229,248 100.00%
Not attended	50,165 5.87%	804,942 94.13%	855,107 100.00%
Cancelled	1,189 0.89%	133,143 99.11%	134,332 100.00%
Total	2,093,331 10.25%	18,331,873 89.75%	20,425,204 100.00%

Notes: Implausible = between 7 p.m. and 7 a.m.; plausible = between 7 a.m. and 7 p.m.

CHAPTER 2. CASEWORKERS' MEETINGS WITH JOB-SEEKERS

Chapter 3

Do Gender-Balanced Committees Select More Women? Evidence from Academic Hiring

Abstract

We study a 2019 reform of academic hiring committees at an Italian university that paired a one-third gender quota with the random assignment of evaluators to committees. The lottery generates quasi-exogenous variation in committee gender composition, which we exploit to identify its effect on candidates' outcomes. Combining committee records, the national registry of academic personnel, and an open bibliographic database, we assemble an original dataset covering 789 hiring competitions held between 2015 and 2024. The reform raised the average share of female evaluators from 22% to 43%. Exploiting the random assignment, we find no evidence that committee gender balance affects a woman's probability of being selected; the estimates rule out sizeable effects but not more modest ones. A difference-in-differences design comparing fields differentially exposed to the reform points, more tentatively, in a positive direction in the most male-dominated fields.

3.1 Introduction

The under-representation of women in academia continues to be a concern for universities and policymakers. While women's participation in academia has increased significantly over time, they remain underrepresented in senior positions and in the most prestigious universities, with strong variation across countries and disciplines (Iaria et al., 2022). In Europe, women account for about 48% of PhD

This chapter is co-authored with Mauro Sylos-Labini (University of Pisa).

graduates, although there are significant disparities persisting between fields and subfield: women represent the majority (over 60%) of Doctoral graduates in Education but are underrepresented (around 22%) in STEM fields such as Information and Communication Technology (ICT) and Engineering (European Commission, 2024). Despite apparent equity at the aggregate PhD level, there is an additional progressive loss of female academics along the path to senior positions, the so-called 'leaky pipeline' (Osterloh and Rost, 2025).

In fact, female representation drops to 40% among associate professors and to below 30% among full professors on average. The decline is considerably steeper in traditionally feminized disciplines: despite entering academia in large numbers, women retain only 38% of full professor positions in the humanities and the Arts, 35% in the social sciences, and 33% in the medical and health sciences. The disparity is even more pronounced in Science and Engineering, where women hold only around 20% of Grade A positions.

More recent cross-country evidence confirms that these patterns are not unique to Europe: studies document persistent gender disparities in academic advancement in the United States, Japan, East-Asia, etc. (Spoon et al., 2023; Nakajima et al., 2023; Cosentino, 2025; Kam and Lee, 2024).

A common expectation, often referred to as the *pipeline theory*, is that once women enter the lower rungs of the academic career ladder, gender gaps at senior levels should gradually narrow as cohorts progress over time. However, empirical evidence shows this convergence is not yet realized, even in cases where the undergraduate and doctoral participation have been increasing (Ceci et al., 2014; Ginther and Kahn, 2006).

Large scale studies dive into this pattern and document gender disparities in both promotion process and career trajectories. Morettini and Tani, 2025 use longitudinal data on European researchers to show that women are less likely to be promoted than men, but those who are promoted often advance faster, indicating strong self-selection and persistence effects. Exploiting the two-stage promotion system introduced in Italian academia, De Paola et al., 2018 show that while no gender differences emerge in the national qualification stage, women face significantly lower probabilities of promotion at the departmental level, particularly when positions are scarce. Complementary evidence by Marini and Meschitti, 2018 indicates that this disadvantage is not confined to scarcity contexts: using multilevel administrative data and detailed controls for scientific productivity and institutional characteristics, they document a persistent and transversal gender gap in promotions, with men retaining higher probabilities of advancement even at parity of merit. At the institutional level, Müller and Palmén, 2025 show that women's mid-career retention depends less on university gender policies and more

3.1. INTRODUCTION

137

on broader structural factors such as childcare systems, job stability, and welfare regimes.

Harris et al., 2025 further develop a model of academic promotions and exits that incorporates statistical discrimination and bias, and empirically show that women tend to be under-represented at higher academic grades and experience longer times to promotion and higher probabilities of exiting academia earlier compared to men.

Looking beyond promotion processes to the entire career trajectory, Huang et al., 2020 document, using longitudinal data on scientific publishing careers across countries and disciplines, that although women's participation in academia has increased over time, gender gaps in total productivity and impact still persist and have even increased. They note that men and women publish at similar rates per year and generate similar impact for the same number of publications per year, which indicates that a lot of the gap is explained by shorter career lengths and higher dropout rates among women rather than lower performance per se.

Beyond this vertical segregation, academia also displays horizontal segregation, since women are much less likely to study and work in STEM fields, which has been linked to potentially large aggregate productivity losses in research due to misallocation of talent (Ahimbisibwe et al., 2025).

The causes behind these patterns are manifold and have been extensively studied, so we briefly summarise the main explanations without claiming exhaustiveness, given the breadth of this literature. Gender differences in promotion rates may reflect differences in productivity, partly shaped by the unequal distribution of household work and childcare (Goldin, 2014; Blau and Kahn, 2017), the direct effect of parenthood (Morgan et al., 2021) or by gendered roles within the household more broadly.

Some authors argue that women may spend an excessive amount of time on tasks that are socially valuable but poorly rewarded in promotion decisions such as service, mentoring, and administrative duties (Babcock et al., 2017). Others highlight that women are less prone to apply for promotions or to engage in competitive selection processes, even when equally qualified (Bosquet et al., 2019; De Paola and Scoppa, 2015), a pattern that has been associated with gender differences in preferences and beliefs about competitive environments (Coffman et al., 2021) and differences in bargaining behavior and perceived social costs of self-advocacy in the labor market (Exley et al., 2020).

Beyond these explanations, existing evidence provides conflicting conclusions on whether female representation in hiring committees improves outcomes for women. One reason is that studies either exploit exogenous variation in committee assignment in the absence of quota reforms (Bagues et al., 2017), or evaluate quota

reforms without quasi-random variation in committee composition (Deschamps, 2024). This chapter combines both sources of variation.

3.1.1 Composition of the committee

Beyond the previous supply side explanations, several contributions link the slow progress of women to the composition and behavior of the committees that decide on hiring and promotions. One strand of the literature investigates whether the presence of women on scientific and hiring committees increases the chances of success of female candidates and improves the quality of evaluations, pointing to three main mechanisms: similarity bias across research fields, gendered professional networks, and taste-based or statistical discrimination. The first mechanism, similarity bias, stems from gender segregation across scientific subfields. Men and women tend to work in different research areas (Dolado et al., 2012). If evaluators overrate the importance of their own type of research, the lack of female evaluators may disadvantage women whose work is concentrated in different niches (Bagues and Perez-Villadoniga, 2013). Recent evidence using NLP-based text embeddings to measure research similarity shows that women have lower maximal similarity to male-dominated committees, suggesting a similarity bias that disadvantages female candidates (Bello et al., 2025).

The second mechanism operates through gendered research networks: coauthoring patterns are not gender neutral in economics (Boschini and Sjögren, 2007), and female PhD students are far more likely to be advised by women than by men, while the opposite holds for male students (Hilmer and Hilmer, 2007). If most evaluators are male, male candidates may be more likely to have prior connections with committee members and to benefit from these ties (Zinovyeva and Bagues, 2015). A related study on recruitment to entry-level research positions in Italy (Checchi et al., 2019) shows that while the presence of at least one woman on the committee can mitigate gender bias at the non-tenured level, hiring outcomes are largely driven by pre-existing connections, an advantage disproportionately enjoyed by male candidates due to denser network ties.

The third mechanism concerns more direct forms of discrimination, whether taste-based or statistical. Survey evidence suggests that a sizable share of men believe that men make better political leaders or business executives, while women are much less likely to endorse such views (OECD, 2025).

The evidence is mixed. A large-scale experiment in Nordic universities finds no systematic bias against women in academic recruitment (Carlsson et al., 2021). A factorial survey of German professors finds slightly more favorable evaluations for female assistant professor candidates, while a parallel experiment in Italy finds

3.1. INTRODUCTION

139

no significant gender differences in evaluations (Gërxhani et al., 2025). Some meta analyses argue that empirical evidence does not provide clear support for systematic discrimination in reviewing, interviewing, or hiring¹ (Williams and Ceci, 2015; Ceci and Williams, 2011), others support the contrary (Coffman et al., 2021).

A related literature examines whether group composition and gender diversity affect decision quality. Experimental and field evidence suggests that more gender-diverse teams can outperform homogeneous ones, although the relationship is often non-linear and context-dependent (Kim and Starks, 2016): applied to academic governance, this perspective implies that a higher share of women on scientific boards could in principle strengthen meritocratic evaluation and improve collective judgment.

However, empirical evidence remains mixed. Even when female evaluators are more supportive of female candidates, they may not exert equal influence in collective deliberations, as women's recommendations and contributions tend to carry less weight in group decision-making. Radbruch and Schiprowski, 2023 show that in hiring committees women's recommendations are systematically less likely to shape the final decision, especially when members disagree and deliberation becomes decisive². In line with this, Karpowitz et al., 2024 show that women exert significantly less influence in male-majority teams and are less likely to be selected as external representatives. However, assigning a female leader or increasing the share of women in the group substantially mitigates this disadvantage, reducing the effective "penalty" faced by female participants by more than half. In line with this, Born et al., 2022 show experimentally that women in male-majority teams are significantly less willing to lead than those in female-majority teams, an effect driven not by performance differences but by lower confidence in relative ability and reduced expectations of electoral support.

Importantly, recent work shows that the effect of committee gender composition is not stable over time but depends on social and organizational contexts. Baron et al., 2024 find that increases in female representation on nomination committees initially reduced female appointments, but turned positive after some years, alongside organizational reforms and rising societal concern for gender diversity. This evidence helps reconcile mixed findings by showing that the impact of female representation in evaluative bodies is contingent on prevailing norms

¹A set of experimental studies focuses directly on discrimination in the evaluation of academic CVs and finds that identical applications or CVs receive lower evaluations when associated with a female name (Steinpreis et al., 1999; Moss-Racusin et al., 2012). These studies are relatively dated compared to the more recent studies.

²This influence gap does not reflect differences in experience or evaluation quality, but emerges during discussion, where women's contributions carry less weight. The effect is strongest when a lone woman sits with two men, implying that adding a single female member may have limited impact on outcomes, whereas gender balance or female majorities substantially reduce the gap.

rather than mechanically positive.

A particularly influential contribution is Bagues et al., 2017, which provides large-scale causal evidence on the effect of committee gender composition by exploiting the quasi-random assignment of evaluators in national academic evaluations in Italy and Spain. Using a natural-experiment setting with an initial random draw and an intention-to-treat design, the authors isolate exogenous variation in committee gender balance. Their results show that increasing female representation in committees does not systematically improve women's promotion chances nor the observable quality of selected candidates, and may even reduce male evaluators' support for female applicants in mixed-gender committees³.

A complementary perspective is offered by Bagues et al., 2026, who shift attention from the composition of committees to the downstream consequences of promotion outcomes. Exploiting the same quasi-random assignment of evaluators in the Spanish habilitación system, they show that exogenously promoting a woman to Associate Professor generates lasting trickle-down effects along the academic pipeline: affected departments have approximately 1.5 additional female Associate Professors fifteen years later, and each female promotion leads to six additional female PhD graduates over the following decade. Strikingly, failing to obtain tenure has asymmetrically lasting consequences for women relative to men, suggesting that early promotion bottlenecks are a key and underappreciated driver of persistent female underrepresentation. Taken together with Bagues et al., 2017, these findings point to a tension at the heart of the literature: while increasing female representation on committees does not straightforwardly improve women's chances, the stakes of each individual promotion outcome are substantially higher for women than for men, making the design of fair and effective evaluation processes all the more consequential.

3.1.2 Gender quotas

While the gender composition of committees may matter, evidence is mixed and positive effects may take time to be achieved. Yet gender quotas have been increasingly adopted, partly because they are the most immediate and enforceable way to ensure balanced representation, and because of their documented successes in other domains concerning top-level representation, notably corporate boards and politics.

³They find that in Italy an additional female commissioner widens the gender gap in success rates by about 1.8 percentage points, driven by male evaluators becoming less favorable toward female candidates, while in Spain they rule out any sizeable aggregate effect, with the exception of promotions to full professor, where female candidates benefit from female representation.

3.1. INTRODUCTION

141

A large share of the literature focuses on board quotas⁴. The broader pattern is confirmed by the meta-analysis of 51 studies across 11 countries by De Acutis et al., 2024, which finds consistent improvements in board quality, but limited vertical spillovers within firms and more frequent negative stock-market reactions. Parallel research on political quotas reaches similar conclusions: quota policies reshape both representation and the profile of elected officials, often without lowering average quality (Baltrunaite et al., 2014).

In Europe, several countries have introduced gender balance requirement in scientific and hiring committees, including at least one third of each gender, and many universities and research institutions have adopted internal guidelines to promote gender-balanced decision-making bodies (European Commission, 2018).

Deschamps, 2024 provides the first direct causal estimate of a legal gender quota for academic hiring committees. Exploiting a French reform that imposed a minimum 40% share of each gender, he uses a difference-in-differences design comparing university-field cells above and below the threshold before the reform. The quota raised the share of women in treated committees from about 30% to 45%, but reduced the probability that a woman was ranked first by about 18 percentage points, with individual-level marginal effects between 8 and 18 percentage points depending on the specification. The reform had no effect on the gender of committee presidents, committee quality, or the gender composition of the applicant pool, and the negative effect is entirely driven by committees headed by men.

A related line of research points to the highly contextual nature of gender bias and acceptance of equality policies in academia. Using harmonized factorial survey experiments in Germany and Italy, Gërxhani et al., 2025 find pronounced cross-country heterogeneity in candidate evaluations: whereas female applicants are evaluated more favorably in Germany, consistent with a policy environment actively promoting gender equality, evaluations in Italy appear largely gender-neutral, reflecting a more traditional normative context and the absence of preferential selection policies. Complementary evidence from social psychology further highlights the role of cultural norms in shaping attitudes toward gender interventions. Esnard and Grangeiro, 2025, comparing female academics in France and Brazil, document stark differences in support for gender quotas: French respondents, embedded in a highly competitive and meritocratic academic culture, are significantly more hostile to quotas and more likely to legitimize existing hierarchies, whereas Brazilian respondents show stronger support for corrective policies

⁴Norway set the benchmark through the early 40% reform (Bertrand et al., 2014). Following Norway, quota policies spread widely across Europe (European Parliament and Council of the European Union, 2022), generating an extensive empirical literature.

and greater solidarity with junior women. Taken together, these findings suggest that both evaluative biases and the legitimacy of gender-equality measures vary systematically across institutional and cultural environments. This heterogeneity implies that evidence from one country cannot be readily extrapolated to another, making it essential to conduct a dedicated analysis of the Italian context.

As anticipated, the existing empirical literature also differs substantially in identification strategies, focusing either on exogenous variation in committee composition without institutional reforms (Bagues et al., 2017) or on quota-based policy changes identified through pre-post designs (Deschamps, 2024). To date, little evidence combines a formal quota introduction with quasi-random variation in committee gender composition at the micro-institutional level.

This chapter contributes to the literature by jointly analyzing a policy-driven quota reform and quasi-random variation in committee assignments within a single university setting, allowing us to disentangle institutional and compositional effects. A further distinguishing feature of the University of Pisa reform is the randomization mechanism itself. To our knowledge, no other Italian university has introduced a comparable lottery-based procedure for committee assignment. This makes the University of Pisa a unique institutional setting in which quasi-random variation in committee gender composition can be credibly exploited for causal inference, and limits the scope for the strategic manipulation of committee composition that may confound observational studies in other contexts. Table 3.1 summarizes the key differences between this chapter and the two closest articles in the literature. Despite focusing on a single institution, this chapter covers 789 hiring committees, exceeding Deschamps, 2024, which includes 455 committees across three French universities.

Table 3.1: Comparison with the two closest studies in the literature.

	Bagues et al. (2017)	Deschamps (2024)	This chapter
Country	Italy & Spain	France (3 uni- versities)	Italy (Univ. Pisa)
Years	1987-2007	2014-2019	2015-2024
Committee type	Evaluation	Hiring	Hiring
Identification strategy	Randomization (ITT)	Pre-post (ranked logit)	Randomization (conditional ATE) + pre- post
Number of committees	100,000	455	789
Avg. women on committee	19%	30% → 45%	22% → 43%

The recruitment procedures in Italy for accessing university careers (full, associate and assistant professor positions) are managed directly by the universities through local competitions, in compliance with universities' autonomy in the framework provided by the Italian law⁵. This implies that the recruitment procedures can differ among institutions.

This chapter is organized as follows. Section 3.2 describes the recruitment procedure at the University of Pisa and the July 2019 gender-equality reform. Section 3.3 presents the data. Section 3.4 exploits the lottery to estimate the effect of committee gender composition on female candidates' probability of winning. Section 3.5 turns to the effect of the reform itself, using a difference-in-differences design across fields differentially exposed to it. Section 3.6 concludes.

3.2 Institutional Background: Recruitment Procedures and the Gender Equality Reform

Academic recruitment at the University of Pisa follows a three-stage process governed by Law 240/2010. A department first requests the opening of a position, approved by the Board of Directors; the dean then issues an official public call specifying the competition scientific sector⁶, the required qualifications, and the evaluation criteria; finally, a committee evaluates the eligible candidates and produces a final report identifying the most suitable one, whom the department council formally appoints.⁷

It should be noted that the structure of committee evaluations varies across positions: while researcher competitions always include a numerical score, evaluations for professor positions are more heterogeneous and do not systematically assign a numerical ranking. In the evaluation report for Full and Associate Professors, an opinion on the candidate is still expressed; however, no numerical score is provided, which in our case makes it more difficult to assign a ranking to candidates for these positions.

For full and associate professor positions, candidates must hold the national qualification⁸ in the relevant competition sector. For assistant professor positions and researchers⁹, a PhD replaces the national qualification requirement, and asso-

⁵Art. 18, comma 1, Legge 240 del 2010; art. 24, comma 3, lettera b) della Legge 240 del 2010

⁶settore scientifico disciplinare, SSD

⁷The call is posted on the university website and on the *Gazzetta Ufficiale*; both the opening of the position and the final appointment require an absolute majority of the department council. For full and associate professor positions, the committee evaluates applicants according to the criteria established by Ministerial Decree n. 344/2011.

⁸Abilitazione Nazionale

⁹Ricercatori a Tempo Determinato, RTD-A and RTD-B

ciate professors with permanent contracts are also eligible as committee members. The procedure for these positions additionally includes a shortlisting phase – retaining between 10% and 20% of applicants, with a minimum of six – followed by a public oral discussion; here the committee assigns a numerical score and the department is bound to appoint the highest-ranked candidate. These procedural differences across position types motivate the inclusion of competition-type fixed effects throughout the empirical analysis.

3.2.1 The Gender Equality Reform of July 2019

Before July 2019, hiring committees included up to three members, preferably full professors, with at least one external to the University of Pisa, and no gender-balance requirement was in place. The reform that came into force on 20 July 2019 introduced two changes: a minimum quota of one third female evaluators, and a lottery in the selection of committee members that makes the quota enforceable.

The composition rules differ by position and are summarized in Table 3.2. For full-professor competitions, committees must include at least five members, two appointed by the university and three drawn by lottery from a gender-contingent shortlist. For associate and researchers competitions, the minimum size is three, one appointed and two drawn from a shortlist of four, again subject to a gender-contingent rule. Relative to the pre-reform regime, this raises the minimum size of full-professor committees from three to five, while associate and assistant committees remain at three. Departments that cannot meet the minimum female representation requirement must formally justify the exception.

Table 3.2: Committee composition rules under the 2019 reform

	Full professor	Associate / Assistant
Minimum committee size	5	3
Members appointed by the university	2	1
Members drawn by lottery	3	2
<i>Shortlist composition</i>		
if appointed member is female	≥ 2 W and ≥ 2 M	≥ 2 W
if appointed member is male	≥ 4 W	≥ 3 W

Notes: W and M denote women and men. For associate and assistant professor competitions the shortlist consists of four candidates; for full professor competitions the gender constraints apply to the shortlist from which the three lottery members are drawn.

Table 3.3 reports the composition realized in practice. Committees adhere closely to these minima: full-professor committees sit typically five, with some

3.2. INSTITUTIONAL BACKGROUND AND THE 2019 REFORM

145

transitional committees below the new minimum, associate and assistant committees at three, and the lottery accounts for roughly two thirds of the seats. The share of women among evaluators, between 0.41 and 0.44 across position types, sits above the one-third floor.

Table 3.3: Effective committee composition by competition type

	N	Members	Appointed	Lottery	Share female
<i>Panel A. Pre-reform</i>					
Full professor	91	3.02	all appointed		0.20
Associate	88	3.00	all appointed		0.19
Researcher	28	3.00	all appointed		0.40
All	207	3.01	all appointed		0.22
<i>Panel B. Post-reform</i>					
Full professor	162	4.77	1.65	3.11	0.40
Associate	66	3.00	1.00	2.00	0.41
Researcher	354	3.01	1.00	2.01	0.44
All	582	3.50	1.18	2.32	0.43

Notes: Average number of seated committee members per competition, split into members appointed by the university and members drawn by lottery, with the share of women among them. Before the reform there was no lottery and all members were appointed. Full-professor committees had five members through 2023 and three from 2024, following the February 2024 revision of the recruitment regulation (D.R. 358/2024), which lowers the post-reform average below five; the lottery still accounts for two of the three seats.

This lottery-based assignment generates **quasi-random variation in the realized gender composition of committees**, which is **the** main source of identification in the empirical analysis. It is worth being precise about what the lottery does and does not randomize. The draw operates within the shortlist: it randomizes which members are selected conditional on the pool, but the pool itself is not drawn from the universe of professors in the sector. The shortlist is assembled by the department and the university according to their internal rules, so the set of names from which the lottery extracts, and in particular which women enter it, remains a discretionary choice. This differs substantively from the centralized mechanisms that preceded, and may again follow, the current regime (Appendix 3.A.1), where the draw is carried out by the Ministry from national lists. The way departments populate the shortlist is therefore itself an object of interest, and we return to it, and to what it implies for identification, in Section 3.4.1.

To place the reform in context, we collected committee-composition rules across Italian universities, reported in Appendix 3.A.2. The Pisa combination, a one-third quota enforced through a lottery, turns out to be rare: only seven of the forty-one universities that responded to our 2023 survey combine a one-third quota

with a random mechanism. This makes Pisa a particularly informative setting for studying this type of intervention.

3.3 Data

We assemble an original dataset from three sources: the publicly available administrative documents of the University of Pisa, the CINECA national registry of academic personnel, and the OpenAlex bibliographic database. The first gives the exact composition of each committee and the list of candidates in each competition; the second the age, academic position, and disciplinary sector of evaluators, and the pre-reform gender composition of each field; the third the research output of every evaluator and candidate, as well as the affiliation.

The competition records are drawn from the calls and their outcomes published on the official website, in the section for selection procedures, in compliance with Italian transparency legislation.¹⁰ They cover competitions opened between 2015 and 2024 for full professor, associate professor, and researcher positions (junior RTD-A and senior RTD-B). Competitions opened before 20 July 2019 form the pre-reform period, those opened from that date the post-reform period. Coverage of the pre-reform years is still being completed, so pre-reform figures are preliminary and 2017 is currently missing; we return to this in Section 3.3.1. For each competition we record the disciplinary sector (SSD), the competition sector and macro-sector, the position type, and whether the call was open only to external candidates or also to internal ones.¹¹ The resulting dataset covers 789 competitions, 207 before and 582 after the reform, with 3,150 candidates.

Bibliometric indicators come from OpenAlex, retrieved through a Python web-scraping procedure.¹² We collect three measures for each evaluator and candidate: the h -index, the number of documents, and the number of citations. For documents and citations we observe both the current value and the value at the time of the competition; for the h -index only the current value is used, since reconstructing it at the time of the competition would require rebuilding the full citation history. Indicators enter the analysis in logs. Name disambiguation restricts matches to authors active in the same disciplinary sector as the competition, using a field-match score between the author's OpenAlex concepts and the sector; profiles with no disciplinary overlap are left missing rather than assigned arbitrarily, and are ab-

¹⁰unipi.it/index.php/concorsi-gare-e-bandi.

¹¹Art. 18, comma 4 and comma 1, Law 240/2010.

¹²OpenAlex (Priem et al., 2022) is an open scholarly knowledge graph covering over 250 million works, freely accessible at openalex.org. Relative to proprietary databases such as Scopus or Web of Science it offers broader coverage of non-English work and of disciplines underrepresented in traditional indexes (Alperin et al., 2024).

sorbed by the missing-value indicators included in all specifications. Institutional affiliation is taken from the competition documents and completed with OpenAlex where missing.

We complement these with CINECA administrative data on the career histories of academic staff in Italian public universities, covering academic position, department, disciplinary sector, age, and gender.¹³ From this source we build the pre-reform female share of each field, which enters the difference-in-differences analysis as a measure of exposure to the reform, and, for Pisa staff, we link the bibliometric measures above through the same OpenAlex procedure.

3.3.1 Limitations of the current dataset

This section states the limitations of the dataset and how we address them. Some reflect the current stage of data collection and will be relaxed as further records become available.

Completeness of the candidate lists. How completely we observe the candidate list depends on the position type. For full and associate professor competitions, we observe the full applicant pool by construction. Candidates submit their applications to the Rector before the call deadline, and the applications are sealed. The committee is appointed only after applications have closed, so it cannot influence who applies. In its first meeting the committee sets the evaluation criteria, and only then opens the sealed list of applicants, transmitted by the Recruitment Office. The pool is therefore fixed before the committee is formed, and the admissibility check is carried out by the administration, not by the committee, which rules out any selection into the sample on the committee's part.

For researcher positions (RTD-A, RTD-B), the procedure includes a preliminary evaluation that admits between 10% and 20% of applicants to the oral discussion, with a minimum of six. In competitions with six or fewer observed candidates, all applicants proceed to the oral stage by rule, so the observed list coincides with the applicant pool. In competitions with more than six observed candidates, the publicly available documents do not let us rule out that the observed list reflects only the candidates admitted to the oral discussion, rather than the full applicant pool.

We address this in two ways. First, we are recovering the complete administrative records for the affected competitions, so their figures will be updated and the present results are preliminary in that respect. Second, for our main specification we assume that the observed candidate list is the complete applicant pool. This

¹³cercauniversita.cineca.it.

holds by construction for professor competitions and for researcher competitions with six or fewer candidates, and may fail only for researcher competitions with more than six. We are explicit that this is an assumption rather than a certainty for that group, and we treat it as a robustness check. Excluding these competitions leaves the interaction statistically indistinguishable from zero, though on a substantially smaller sample whose minimum detectable effect is large enough that the exercise is only weakly informative (Table A5). This reflects the data currently available; recovering the complete oral-stage records for the affected competitions, which is ongoing, will allow us to keep them in the full sample rather than exclude them.

A single institution. Our identification rests on the lottery-based assignment of committee members. Very few Italian universities pair a binding one-third quota with a random draw (Appendix 3.A.2), and Pisa is the one for which we have complete administrative records of committee composition and outcomes. This is what makes the setting valuable, but it also means our evidence comes from a single institution, which limits external validity. Whether the estimated effects extend to universities with different disciplinary mixes, recruitment cultures, or gender norms cannot be established within our data. The single-institution design also constrains the difference-in-differences analysis: a fully credible pre-post comparison would benefit from a control group of comparable universities that did not adopt the reform. Such a group is hard to identify, since we do not observe with certainty which other institutions introduced analogous measures and when; we therefore treat any cross-university comparison as descriptive and rely on within-Pisa variation for identification. Extending the analysis to a second institution with a comparable mechanism is a natural direction for future work.

An incomplete pre-reform period. The data span 2015–2024, but coverage of the pre-reform years is uneven: 2017 is missing and 2018 is sparsely populated. The pre-reform window available for the difference-in-differences analysis is therefore short and irregular, which prevents a fully convincing test of parallel pre-trends. We are explicit about this throughout: the pre-post results are exploratory rather than clean causal estimates of the reform, and the post-reform lottery analysis, which does not rely on a pre-period, remains the primary source of identification. Completing the pre-reform records is part of the ongoing data collection.

Features of the evaluation being coded. Two elements of the evaluation that the literature has shown to matter are still being coded from the committee minutes. First, the gender of the committee president, recorded in the minutes of the first

meeting: we have coded it for a subset of competitions and are extending it to the rest. Deschamps, 2024 finds that the adverse effect of the quota is driven entirely by male-headed committees, so completing this variable will let us test that channel. Second, researcher competitions assign a numerical score and hence a full ranking of candidates, which we have not yet used as an outcome; exploiting the rank rather than the binary win indicator, as in Deschamps, 2024, would draw identifying variation from every pairwise comparison and could sharpen the estimates for these positions. For full and associate professor competitions this is not possible, as the evaluation expresses a reasoned opinion without a cardinal score. Neither addition affects the validity of the current estimates. Additionally, we are constructing measures of pre-existing ties between candidates and evaluators, which the literature identifies as a channel through which committee composition operates (Zinovyeva and Bagues, 2015; Checchi et al., 2019). From the bibliographic data we can flag whether a candidate has co-authored with any committee member, and from affiliations whether at least one member belongs to the candidate's own university. These links let us test whether the effect of committee composition works through connections rather than through gender as such, and to check that the lottery does not systematically place candidates before better-connected committees.


3.3.2 Summary Statistics

The first thing to establish is whether the reform was actually followed. Table 3.3 shows that it was: after 2019 committees comply closely with the new composition rules. The average share of women among evaluators roughly doubled, from 0.22 before the reform to 0.43 after, moving committees from well below the one-third floor to comfortably above it. The shift is visible across the whole disciplinary distribution (Figure 3.1): all fourteen CUN areas raise their female committee share after the reform, and all but one clear the one-third threshold, the exception being Earth Sciences, a small area that also started from the lowest base. The largest increases are concentrated in the fields that were most male-dominated to begin with, Physics (0.06 to 0.41), Civil Engineering (0.17 to 0.52), and Medicine (0.15 to 0.40), which is what the lottery mechanism should produce: the quota binds hardest, and moves composition most, precisely where the pre-reform pool was furthest from parity. Figure 3.2 shows the same pattern for the lottery-expected share, indicating that the increase reflects the composition of the shortlists, not a particular realization of the draw.

This change in committee composition is not accompanied by a decline in observable committee quality. Post-reform committees are younger on average, with

mean evaluator age falling from about 60 to 57 years, and their measured bibliometric output is no lower than before.¹⁴ The candidate side moves little: the share of female candidates rises only modestly, from 0.27 to 0.31, so the reform changed who evaluates rather than who applies. Table 3.4 reports the same breakdown by CUN area in numbers, giving for each field the count of competitions and the share of women among those hired, on committees, and among applicants. Table 3.5 collects these pre- and post-reform means, and Tables A1 and A2 report them by year and by position type.

Figure 3.1: Average share of female committee members by discipline, before and after the reform (actual).




Notes: Each point is the average realized female share of committees in a CUN discipline, before (blue) and after (orange) the reform. The dashed line marks one third. 2017 is excluded.

¹⁴Average committee documents and citations are somewhat higher post-reform; we do not read this as evidence that individual evaluators are stronger, since bibliometrics are measured up to the competition year and post-reform competitions are more recent, so they accumulate output in a period of generally higher publication volume.

3.3. DATA

151

Figure 3.2: Average share of female committee members by discipline, before and after the reform (expected).



Notes: As in Figure 3.1, using the lottery-expected female share instead of the realized one.

Table 3.4: Descriptive statistics on recruitment by CUN area

CUN Area	N	Pre	Post	Sh. hired	Sh. comm.	Sh. appl.
Mathematics and Computer Science	66	19	47	0.23	0.32	0.18
Physics	49	15	34	0.31	0.31	0.15
Chemistry	47	10	37	0.45	0.47	0.38
Earth Sciences	21	7	14	0.24	0.17	0.20
Biology	59	15	44	0.61	0.47	0.59
Medicine	121	32	89	0.35	0.33	0.39
Agricultural Sciences	57	8	49	0.47	0.40	0.49
Civil Engineering	18	4	14	0.33	0.44	0.34
Industrial Engineering	126	33	93	0.19	0.31	0.16
Architecture	63	15	48	0.41	0.49	0.45
History	34	9	25	0.50	0.38	0.39
Law	43	15	28	0.25	0.37	0.30
Economics	68	24	44	0.42	0.39	0.36
Political Science	17	1	16	0.29	0.48	0.22
Total	789	207	582	0.35	0.37	0.30

Table 3.5: Competitions before and after the reform

Variable	Pre	Post
N. competitions	207	582
Share F. committee	0.22	0.43
Share F. candidates	0.27	0.31
Share F. winners	0.32	0.36
N. candidates	3.79	4.07
Cand. documents (at comp.)	93.83	113.16
Cand. citations (at comp.)	3112.65	4638.71
Cand. documents (current)	170.43	166.48
Cand. citations (current)	4835.66	5748.61
Cand. H-index (current)	24.79	23.64
Avg. documents committee (at comp.)	171.81	194.40
Avg. citations committee (at comp.)	4679.92	5530.72
Avg. age committee	60.24	57.38
Share Pisa committee	0.42	0.21
N. candidates (obs.)	784	2366

3.4 Does Committee Gender Composition Affect the Probability of Winning for Female Candidates?

3.4.1 Empirical Strategy

We now turn to the central question: does the gender composition of the hiring committee affect the probability of success of female applicants? The reform introduced two sources of variation. We estimate the following linear probability model at the candidate-by-competition level:

$$\begin{aligned}
 Y_{ic} = & \beta_0 + \beta_1 \text{Female}_i + \beta_2 \text{ShareFemale}_c^{\text{actual}} \\
 & + \beta_3 \left(\text{Female}_i \times \text{ShareFemale}_c^{\text{actual}} \right) + \beta_4 \text{ShareFemale}_c^{\text{expected}} \\
 & + \beta_5 \left(\text{Female}_i \times \text{ShareFemale}_c^{\text{expected}} \right) + X'_{ic} \delta + \mu_s + \tau_t + \varepsilon_{ic}
 \end{aligned} \tag{3.1}$$

where Y_{ic} equals one if candidate i wins competition c ; Female_i indicates a female candidate; $\text{ShareFemale}_c^{\text{actual}}$ is the realized female share of the committee, fixed by the lottery draw; and $\text{ShareFemale}_c^{\text{expected}}$ is the female share the lottery is expected to produce, holding the shortlist and the directly appointed members fixed. The expected share is built seat by seat from the draw probabilities: a directly appointed member contributes one if she is a woman and zero otherwise, while each woman eligible for a lottery seat contributes her probability of being drawn (1/2 or 1/3, for instance, depending on how many seats are drawn from the shortlist and how many candidates are eligible). Summing these contributions and

3.4. DOES COMMITTEE COMPOSITION AFFECT WOMEN'S SUCCESS? 153

dividing by committee size gives $\text{ShareFemale}_c^{\text{expected}}$. The coefficient of primary interest is β_3 , which captures whether female candidates benefit differentially from a more female committee, holding the pre-lottery composition fixed.

The vector X_{ic} collects candidate-level and committee-level controls: the number of candidates in the competition, an indicator for affiliation with the University of Pisa, the candidate's own productivity at the time of the competition (log documents and log citations), the average age and bibliometrics of the committee (average documents, citations and h -index), the actual and expected share of Pisa-affiliated committee members, and indicators for missing values. μ_s and τ_t denote scientific-sector (SSD) and year fixed effects. In the robustness analysis we enter candidate productivity as within-competition percentile ranks rather than logs; the results are unchanged (Table A12).

The identifying content of equation (3.1) comes from including the actual and the expected share together. The shortlist from which the lottery draws is assembled by the department, so $\text{ShareFemale}_c^{\text{expected}}$ could be a source of endogeneity: departments that populate more female shortlists may differ systematically in ways correlated with candidate outcomes. By construction the expected share is a deterministic function of the shortlist and the appointed members, computed from the draw rules and the set of eligible evaluators. The causal interpretation of β_2 and β_3 relies on the assignment being random conditional on the shortlist: the department's discretion enters through which names populate the shortlist, and the lottery then draws within it. Conditional on $\text{ShareFemale}_c^{\text{expected}}$, the realized share varies only through the draw, and is therefore orthogonal to anything fixed before the lottery is run. The expected share absorbs the part of committee composition the department selects through the shortlist, leaving the draw-induced variation to identify β_2 and β_3 . We verify this directly: the randomization checks in Section 3.4.2 show that, conditional on the expected share, the actual share is uncorrelated with predetermined candidate and competition characteristics. We do not read β_4 as the effect of the reform: within the post-reform sample it captures the cross-sectional association between an endogenous shortlist composition and the outcome. The effect of the reform itself is estimated in Section 3.5.

Our preferred specification includes year and SSD fixed effects. Because the realized composition is quasi-random conditional on the expected share, the fixed effects serve precision rather than identification, and Table 3.7 reports the full progression. The interaction β_3 is positive and statistically insignificant in every column, although its magnitude varies across specifications. Procedural rules differ across position types, so we also add competition-type fixed effects (PO / PA / RIC / RTT); specifications that include them yield the same conclusion (Table A13).

Standard errors are clustered at the SSD level. The individual competition

would be a natural unit, since candidates in the same competition face the same committee and compete against one another, but clustering at the SSD level additionally captures correlation across competitions in the same field, induced by shared evaluation norms and field-specific hiring patterns. The choice does not drive inference: clustering at the competition level instead leaves the standard error on β_3 essentially unchanged (Table A14).

3.4.2 Main results

We begin with the randomization check in Table 3.6. Each column regresses a predetermined candidate or competition characteristic on the actual and the expected female share, holding year and SSD fixed effects constant. If the draw is as good as random within the shortlist, the actual share should be unrelated to anything fixed before the lottery, once we condition on the expected share. This is what we find. The coefficient on the actual share is small and insignificant across the board, with a single marginally significant estimate (the rank of documents, at the ten percent level) of the kind nine independent tests produce by chance. The expected share behaves differently, as it should. It is a deterministic function of the shortlist, and it correlates with committee h -index and average age, the dimensions along which departments assemble the pool. The contrast is the point: the part of committee composition the department selects is correlated with predetermined characteristics, while the draw-induced part we exploit is not.

Table 3.6: Randomization check: predetermined characteristics on actual and expected female share.

	Log docs	Rank docs	Log cites	Rank cites	Gender	Pisa	N. candidates	H-index comm.	Age comm.
Share female (actual)	−0.077 (0.218)	−0.036* (0.021)	−0.227 (0.308)	−0.031 (0.021)	0.025 (0.053)	0.051 (0.061)	−2.288 (1.846)	0.093 (0.080)	0.830 (1.453)
Share female (expected)	0.202 (0.744)	0.098 (0.094)	0.545 (1.249)	0.113 (0.094)	0.008 (0.092)	−0.025 (0.089)	6.098 (4.379)	−0.454*** (0.161)	−5.632** (2.354)
Num.Obs.	2240	2240	2240	2240	2240	2240	2240	2240	2240
R2	0.381	0.078	0.487	0.076	0.259	0.215	0.710	0.955	0.626

Notes: Each column regresses a predetermined candidate or competition characteristic on the actual and expected female share of the committee, with SSD and year fixed effects. SE clustered at the SSD level in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.7 reports the main specification, adding fixed effects progressively across columns. The coefficient of interest is the interaction between female candidacy and the actual female share, β_3 . It is positive in every column and never approaches significance, with point estimates between 0.019 and 0.051. We find no evidence that a more female committee raises a woman's probability of winning relative to a man's. The main effect of female candidacy is close to zero throughout, and the interaction with the expected share is equally imprecise.

3.4. DOES COMMITTEE COMPOSITION AFFECT WOMEN'S SUCCESS?

155

Table 3.7: Effect of committee gender composition on the probability of winning (contestable sample, $n > 1$).

	(1) No FE	(2) Year	(3) CUN	(4) SSD	(5) Year + CUN	(6) Year + SSD
Female candidate	0.007 (0.052)	0.008 (0.052)	-0.004 (0.052)	-0.008 (0.061)	-0.001 (0.053)	-0.009 (0.062)
Share female (actual)	-0.043 (0.029)	-0.024 (0.031)	-0.052* (0.027)	-0.094*** (0.031)	-0.031 (0.029)	-0.067* (0.036)
Female × Share female (actual)	0.038 (0.072)	0.019 (0.075)	0.046 (0.073)	0.051 (0.080)	0.028 (0.076)	0.037 (0.082)
Share female (expected)	0.057 (0.046)	0.028 (0.046)	0.055 (0.047)	0.083* (0.047)	0.026 (0.048)	0.065 (0.053)
Female × Share female (expected)	-0.006 (0.120)	0.010 (0.121)	-0.007 (0.123)	0.010 (0.141)	0.004 (0.124)	0.025 (0.144)
Number of candidates	-0.009** (0.004)	-0.010*** (0.003)	-0.009** (0.004)	-0.008** (0.004)	-0.009** (0.003)	-0.008** (0.003)
Log docs (until competition)	0.040*** (0.013)	0.041*** (0.014)	0.038** (0.015)	0.052*** (0.016)	0.038** (0.015)	0.052*** (0.017)
Log cites (until competition)	-0.017** (0.007)	-0.017** (0.007)	-0.016* (0.008)	-0.019** (0.009)	-0.015* (0.008)	-0.019** (0.009)
Num.Obs.	2240	2240	2240	2240	2240	2240
R2	0.132	0.135	0.136	0.173	0.139	0.175

Notes: LPM at the candidate-by-competition level; the outcome equals one if the candidate wins. Columns differ in the set of fixed effects. SE clustered at the SSD level in parentheses. Controls included but not shown: Pisa candidate dummy, committee average age and bibliometrics, actual and expected committee Pisa share, missing-value indicators. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The coefficient on the actual share on its own, β_2 , is negative and marginally significant in the preferred specification (-0.067 , column 6). We do not lean on it: it captures the association for male candidates, it is small, and it moves across specifications, while the quantity the design is built to identify is the differential effect for women, β_3 , which is a precise zero. The controls behave as expected. Candidate productivity matters, with the log number of documents at the time of the competition positive and significant, while the number of candidates enters negatively and mechanically, since one candidate wins. The negative coefficient on log citations reflects the strong correlation between the two bibliometric measures rather than a genuine negative effect of citations on winning. Affiliation with the University of Pisa is the one large predictor of success, worth roughly 32 percentage points (Table A6); it does not bear on the committee-composition channel of interest.

The same null holds when the analysis is run at the competition level, with the gender of the winner as the outcome and one independent observation per competition (Appendix 3.A.4).

3.4.3 Statistical power and minimum detectable effects

The estimates are precise in design but imprecise in magnitude. The lottery delivers exogenous variation in committee composition, so the coefficient of interest is cleanly identified. What a single university over a limited period does not provide is the precision to pin down its size. Table A16 reports the minimum detectable effect (MDE) at 80% power and $\alpha = 0.05$.¹⁵ It is about 23 percentage points in the preferred specification and ranges from 20 to 23 across specifications. The full sample gives essentially the same, 21 to 23, so the choice of sample does not drive the result.

The reason is that the design has little identifying variation to work with, and no amount of candidates fixes it. The treatment varies at the committee level, since every candidate in a competition faces the same committee, so what identifies the effect is variation across committees, not across candidates. The contestable sample has 2,240 candidates but only 456 committees, and it is the second number that governs precision: adding more candidates to existing competitions does nothing, only more competitions would help. On top of this, the draw itself moves composition little. The realized female share departs from its lottery expectation by 0.13 on average, so the quasi-random shock we exploit is small to begin with. The granular fixed effects that make the comparison clean then absorb part of even this: moving to year and SSD fixed effects removes about 15% of the residual variation in the treatment while cutting the residual variance of the outcome by only 3%, so the standard error rises rather than falls. The sector dimension binds hardest, since 54% of sectors contain a single competition and contribute no within-sector variation once the fixed effects are partialled out. None of this is a flaw in the design; it is the price of identifying the effect from a clean lottery in a single university over a few years.

These considerations delimit what the design can say. The estimates rule out large marginal effects of committee composition but cannot exclude smaller ones; halving the MDE would require roughly four times as many competitions, attainable only through more years of data or additional universities that also sortition their committees. The value of the exercise is therefore the clean identification it provides at the level of individual hiring, where random assignment of committees is rare, rather than a precise magnitude. A magnitude comparison with the literature is not attempted here. The closest reference, Deschamps (2024), estimates the effect of the reform itself rather than the marginal effect of committee composition; the pre-post analysis of Section 3.5 comes nearer to that parameter, but it is exploratory, and we treat its contribution as the stability of the sign rather than as

¹⁵The MDE is defined as $(z_{0.975} + z_{0.80}) \times \text{cSE} \times 100 = 2.80 \times \text{cSE} \times 100$, where cSE is the clustered standard error of the coefficient of interest.

3.5. DID THE REFORM IMPROVE OUTCOMES FOR WOMEN?

157

an estimate to be compared in levels.

Additional robustness. The null is stable across sample definitions and functional form. Extending the estimation to the full sample (Table A7) and restricting it to competitions with at least one candidate of each gender (Table A8) leaves the interaction of interest indistinguishable from zero. Replacing the linear probability model with a logit (Table A9), and building the specification up through a progressive set of OLS controls on both the contestable and the full sample (Tables A10 and A11), yields the same conclusion. Splitting the sample into STEM and non-STEM fields (Table A15) leaves the interaction a precise zero in both, so the overall null does not mask an effect concentrated in one broad group of disciplines.

3.5 Did the Reform Improve Outcomes for Women?

3.5.1 A continuous measure of treatment intensity

The lottery design of Section 3.4 identifies the marginal effect of committee composition within the post-reform regime. It cannot speak to the effect of the reform itself, because the reform applies to every competition after July 2019 at once: with no untreated post-reform committees, there is no contrast to exploit. To recover the effect of the reform we turn to pre-post variation, and identify it from the fact that the reform did not bite equally everywhere. A one-third quota moves committee composition a great deal in a field that started far below parity and hardly at all in a field already near it. The differential change across fields, interacted with the pre-post indicator and with candidate gender, is the source of identification for the difference-in-differences: it compares the change in female candidates' outcomes, before versus after, across fields differentially exposed to the reform.

The treatment therefore does not switch on and off; it bites more where a field started further from gender parity. We need a measure of how male-dominated each field was before the reform. Male dominance is not directly observable, so we proxy it with the pre-reform share of women in the relevant population, coded as $1 - \text{share female}$, so that higher values denote more male-dominated fields and, mechanically, a larger reform-induced shift in committee composition.

Which population provides the best proxy is not innocuous, and we consider three: the share of women among pre-reform committee members, among national faculty in the field, and among eligible Pisa faculty. We prefer the last, for three reasons that we state up front, precisely so that the choice is not driven by which measure yields the largest coefficient. First, it is external to the outcome. The dependent variable is whether a candidate wins, and the channel we study is

the gender composition of committees; the eligible-faculty pool is the stock from which committee members are drawn, not the committees themselves, so it carries no mechanical link to the object whose variation we are trying to explain. Building the measure from pre-reform committees reintroduces exactly this link. Second, it is Pisa-specific: it reflects the local pool the lottery actually draws from, which the national share approximates only coarsely. Third, it is precisely measured, computed on the full population of faculty in each field, whereas a committee-based measure rests on the thin and incomplete pre-period (207 competitions, with 2017 missing) and is estimated from few committees per field. The three measures are reported by sector in Appendix 3.A.5 (Table A23) and are strongly correlated across fields (Table A24, pairwise correlations between 0.75 and 0.82), so the choice among them affects precision rather than the sign of the results. We report all three throughout and treat the national and committee-based versions as robustness.

We use this proxy as a continuous variable rather than as a binary treated/control split. A natural alternative, in the spirit of Deschamps, 2024, would dichotomize fields, assigning treated status only to those whose committees fell below a female-share threshold before the reform. With only fourteen CUN sectors, such a split discards the within-group spread in male dominance and collapses identification onto a single contrast between two coarse bins, leaving few treated cells and a correspondingly large loss of power; the choice of threshold would itself add a researcher degree of freedom. The continuous measure retains the full cross-sectional variation in pre-reform male dominance across fields, which is the only dimension along which the reform's differential bite can be identified, and avoids committing to an arbitrary cutoff. Before turning to candidate outcomes, we verify that the reform did what the design assumes, in two steps. The first is descriptive: did the reform move committee composition at all? The second is the one the difference-in-differences actually rests on: did it move it more in the fields that started further from parity? An average shift is consistent with the quota having been implemented; only a differential shift makes the cross-field comparison of Section 3.5.2 informative.

We begin by regressing the realized female share of the committee on the post-reform indicator,

$$\text{ShareFemale}_c = \alpha + \beta_1 \text{Post}_c + \delta \mathbf{X}_c + \mu_s + \mu_k + \varepsilon_c, \quad (3.2)$$

where ShareFemale_c is the female share on the committee of competition c , Post_c equals one after 20 July 2019, $\mathbf{X}_c$ collects committee-level controls (average age, citations, documents, and missing-value indicators), and μ_s and μ_k are SSD (or CUN) and competition-type fixed effects. Year fixed effects are deliberately ex-

3.5. DID THE REFORM IMPROVE OUTCOMES FOR WOMEN?

159

cluded: since $Post_c$ switches at a single date, they would absorb all identifying variation and leave β_1 unidentified. Standard errors are clustered at the SSD level.

Table 3.8 shows that the reform bit: the share of female evaluators rose by about 20 percentage points after July 2019, stable across specifications and unaffected by the committee controls. This establishes compliance and nothing more. It says the quota was implemented; it does not say it was implemented unevenly, which is the property the difference-in-differences needs.

We therefore add the interaction between the post-reform indicator and treatment intensity,

$$\text{ShareFemale}_c = \alpha + \beta_1 \text{Post}_c + \beta_2 (\text{Post}_c \times \text{Intensity}_s) + \delta \mathbf{X}_c + \mu_s + \mu_k + \varepsilon_c, \quad (3.3)$$

where Intensity_s is the treatment intensity of field s defined above. Intensity is centered at the contestable-sample mean, so β_1 continues to be read at the average field; being constant within field, its main effect is absorbed by μ_s and does not appear in the columns with fixed effects. The coefficient of interest is β_2 . Since intensity is coded as one minus the pre-reform female share, a positive β_2 means the reform raised the female committee share by more in the fields that started further from parity. This is the same interaction that enters equation (3.4) below, measured there on candidates' outcomes and here on the committee composition through which the reform is meant to operate. Showing it on composition first turns the differential exposure of fields from an assumption of the design into a measured quantity.

The choice of proxy matters more here than elsewhere. The outcome of equation (3.3) is itself the female committee share, so using the pre-reform committee share as the moderator would place the same object on both sides of the regression. With a thin pre-period and few committees per field, a field whose pre-reform share happens to be low by chance would mechanically display a larger subsequent increase, biasing β_2 upward. The eligible-faculty measure is computed on the population of faculty rather than on the committees that form the outcome, and is not exposed to this concern. We therefore use it here, and do not report the committee-based counterpart of this regression.

Table 3.9 reports it. Because intensity is centered, the coefficient on Post continues to be read at the average field, and it is accordingly unchanged. The interaction term is positive in every column. It is estimated precisely under CUN fixed effects (0.473 and 0.480, significant at the one percent level) and less so under SSD fixed effects (0.304 and 0.336), where the finer fixed effects drop the 101 competitions that are the only one in their sector and thereby remove part of the cross-field variation that identifies the interaction. Taken at the estimates of column (5), the

implied rise in the female committee share runs from about 7 percentage points in the least male-dominated field to about 26 in the most. This differential is the variation the difference-in-differences exploits. It is not itself the effect on candidates, which we estimate next.

Table 3.8: First stage: effect of the reform on the female committee share.

	(1) No FE	(2) SSD+Comp	(3) CUN+Comp	(4) SSD+ctrl	(5) CUN+ctrl
Post reform	0.203*** (0.019)	0.209*** (0.023)	0.174*** (0.022)	0.218*** (0.021)	0.180*** (0.021)
Avg. age				0.000 (0.003)	−0.002 (0.002)
Avg. cites (log)				−0.006 (0.014)	−0.004 (0.013)
Avg. docs (log)				−0.052 (0.032)	−0.030 (0.031)
Num.Obs.	789	688	789	688	789
R2	0.125	0.491	0.202	0.499	0.210

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Committee level. The outcome is the realized female share of the committee. SE clustered at the SSD level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.5. DID THE REFORM IMPROVE OUTCOMES FOR WOMEN?

161

Table 3.9: First stage: effect of the reform on the female committee share, average and differential.

	(1) No FE	(2) SSD+Comp	(3) CUN+Comp	(4) SSD+ctrl	(5) CUN+ctrl
Post reform	0.203*** (0.019)	0.209*** (0.023)	0.178*** (0.022)	0.220*** (0.022)	0.184*** (0.021)
Treatment intensity	-0.780*** (0.133)				
Post × Intensity	0.494*** (0.154)	0.304* (0.184)	0.473*** (0.174)	0.336* (0.183)	0.480*** (0.174)
Avg. age				0.000 (0.003)	-0.002 (0.002)
Avg. cites (log)				-0.005 (0.014)	-0.003 (0.013)
Avg. docs (log)				-0.057* (0.031)	-0.032 (0.030)
Num.Obs.	789	688	789	688	789
R2	0.176	0.494	0.212	0.503	0.220

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Committee level. The outcome is the realized female share of the committee. Column (1) reports the average effect of the reform; the remaining columns add field and competition-type fixed effects and, in columns (4) and (5), committee controls (average age, citations, documents, and missing-value indicators). Treatment intensity is one minus the pre-reform female share among Pisa faculty in the sector, centered at the contestable-sample mean, so that the coefficient on *Post* is read at the average field. Higher values denote more male-dominated fields, so a positive interaction means the reform moved committee composition more where the field started further from parity. Being constant within field, the main effect of intensity is absorbed by the fixed effects in columns (2)–(5). SE clustered at the SSD level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.5.2 The reform's differential effect across fields

We estimate the following linear probability model at the candidate-by-competition level:

$$\begin{aligned}
 Y_{ic} = & \alpha + \beta_1 \text{Female}_i + \beta_2 \text{Post}_c + \beta_3 \text{Intensity}_s \\
 & + \beta_4 (\text{Post}_c \times \text{Female}_i) + \beta_5 (\text{Post}_c \times \text{Intensity}_s) + \beta_6 (\text{Female}_i \times \text{Intensity}_s) \\
 & + \beta_7 (\text{Post}_c \times \text{Female}_i \times \text{Intensity}_s) + \mathbf{X}'_{ic} \delta + \mu_s + \mu_k + \varepsilon_{ic}
 \end{aligned}
 \tag{3.4}$$

where Y_{ic} equals one if candidate i wins competition c ; Female_i marks a female candidate; Post_c equals one after 20 July 2019; and Intensity_s is the treatment intensity of field s , one minus the pre-reform female share among Pisa faculty, centered at the sample mean so that the lower-order coefficients are read at the average field. The coefficient of interest is β_7 : whether the reform raised female candidates' winning probability by more in the fields it changed the most. Controls, fixed effects, and clustering are as in the post-reform analysis.

Table 3.10 reports the estimate on the contestable sample, adding fixed effects across columns. The triple interaction is positive throughout and grows as the fixed effects tighten, from 0.285 without fixed effects to 0.490 with CUN and 0.604 with SSD fixed effects, significant at the five percent level in the preferred SSD specification. At the average field the reform leaves female candidates' outcomes unchanged, as the Female $\times$ Post term is small and insignificant; the effect appears only as fields become more male-dominated. Evaluated at the intensity of the most male-dominated fields this is a sizeable differential, but we do not read its magnitude as a clean causal estimate, for the reasons given below.

The sign is stable. It holds on the full and both-genders samples (Tables A25 and A26), under the national and Pisa-committee intensity measures (Tables A27 and A28), and when the potentially truncated researcher competitions are excluded, both under the preferred intensity (Tables A29 and A30) and under the alternative measures (Tables A31 and A32). Clustering at the competition rather than at the SSD level leaves the coefficient unchanged and raises its standard error from 0.275 to 0.322, so that it remains significant at the ten percent level (Table A34). A leave-one-out exercise dropping one CUN sector at a time (Figure A2) leaves the coefficient positive in all fourteen cases, ranging from 0.404 when Architecture is excluded to 0.796 when Industrial Engineering is, so no single discipline drives the result. The confidence intervals are wide, and two of the fourteen include zero.

The balance test for the triple interaction is less clean than for the post-reform lottery. The triple coefficient is insignificant for candidate bibliometrics but marginally significant for Pisa affiliation (-0.627 , $p < 0.1$; Table A33). This is consistent with the exploratory nature of the difference-in-differences: with an incomplete pre-period and treatment intensity defined at the sector level, we cannot rule out that the triple interaction is correlated with predetermined characteristics, and we read the pre-post results as suggestive rather than as clean causal estimates.

We do not report a competition-level version of the difference-in-differences. Collapsing to one observation per competition leaves too few units to estimate the triple interaction with any precision, particularly once fixed effects absorb the sector-level intensity that identifies it, and the resulting estimates are too imprecise to be informative. The candidate-level analysis above is our preferred design for this reason. Recovering a competition-level specification that handles the loss of power appropriately, for instance by modelling the ranking of candidates rather than the binary winner, is left for future work.

We read this result as suggestive rather than conclusive, for the two reasons already noted: the pre-reform window is short and uneven, so parallel pre-trends cannot be convincingly tested, and the share of potentially truncated competitions

3.6. CONCLUSION

163

is itself correlated with the reform. The lottery analysis of Section 3.4 remains the cleaner design; the difference-in-differences complements it by pointing, tentatively, to where in the disciplinary distribution the reform did the most.

Table 3.10: Triple difference on the probability of winning (contestable sample, $n > 1$).

	No FE	CUN FE	SSD FE
Female	0.042 (0.034)	0.005 (0.033)	0.011 (0.034)
Post reform	−0.009 (0.015)	0.006 (0.016)	−0.005 (0.016)
Treatment intensity	0.010 (0.126)		
Female × Post	0.013 (0.041)	0.014 (0.040)	0.012 (0.041)
Female × Intensity	0.064 (0.240)	−0.186 (0.224)	−0.188 (0.244)
Post × Intensity	−0.263* (0.137)	−0.051 (0.117)	−0.072 (0.122)
Female × Post × Intensity	0.285 (0.285)	0.490* (0.264)	0.604** (0.275)
Num.Obs.	2965	2965	2965
R2	0.007	0.148	0.181

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: LPM at the candidate-by-competition level. Treatment intensity is one minus the pre-reform female share among Pisa faculty in the sector, centered at the contestable-sample mean, so the lower-order coefficients are read at the average sector. The coefficient of interest is *Female × Post × Intensity*. SE clustered at the SSD level. Controls included but not shown: candidate and committee covariates and missing-value indicators. **Exploratory:** the pre-reform period is incomplete, so parallel pre-trends cannot be tested. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.6 Conclusion

This chapter asked whether gender-balanced hiring committees select more women, using a 2019 reform at the University of Pisa that paired a one-third gender quota with the random assignment of evaluators. The two designs we bring to the question return, on their face, opposite signs, and it is worth being explicit about why they do not conflict.

The lottery design isolates the marginal effect of committee composition within the post-reform regime: holding the reform fixed, it asks whether a committee

that the draw made more female selects women at a higher rate. It does not. The difference-in-differences design asks a different question: not what one additional woman on the committee does, but what the reform as a whole did, comparing outcomes before and after across fields differentially exposed to it. Here the estimates point, tentatively, in a positive direction in the most male-dominated fields. The two parameters are not on the same scale, as we note in Section 3.4.3: the first is the effect of moving composition at the margin, the second the effect of the whole package, quota and lottery together, in the fields where it bit hardest.

Neither design is conclusive on its own, and they are not conclusive in the same way. The lottery estimate is cleanly identified but imprecise: it rules out large marginal effects of committee composition but cannot exclude smaller ones, so the null is an absence of evidence of a large effect, not evidence of no effect. The difference-in-differences estimate is the opposite, suggestive in sign but weakly identified, resting on a short and uneven pre-reform window that does not allow a convincing test of parallel trends. What the two share is their direction of implication: neither produces the sizeable pro-female effect that a mechanical reading of gender balance would predict, and the only movement we detect is confined, tentatively, to the fields that started furthest from parity.

The contribution of this chapter is therefore as much in the design as in the point estimates. Bringing together, within a single institution, the random assignment of evaluators exploited by Bagues et al., 2017 and the quota-based pre-post variation studied by Deschamps, 2024 lets us ask both what committee composition does at the margin and what the reform does as a whole, two questions the existing literature answers separately. The estimates we report are a first pass: the lottery analysis already delivers causal evidence, at the level of individual hiring, that committee gender composition on its own does not determine whether a woman is selected, while the pre-post comparison awaits the completion of the pre-reform records and, ideally, a second institution with a comparable mechanism before its magnitude can be read.

The limitations that bound these estimates are the ones we have been explicit about throughout: a single institution, a lottery unique to it, and a pre-reform period still being completed. They also mark the natural next steps. Extending the analysis to a second university with a comparable mechanism, and completing the pre-reform records, would turn the difference-in-differences from suggestive into a test, and would let the reform's overall effect be read as a magnitude rather than a direction.

References

- Ahimbisibwe, Isaac, Adam Altmejd, Georgy Artemov, Andrés Barrios Fernández, Aspasia Bizopoulou, Martti Kaila, Jin-Tan Liu, Rigissa Megalokonomou, Jose Montalban, Christopher Neilson, et al. (2025). "Pipeline vs. choice: the global gender gap in STEM applications". In: *IZA Discussion Paper*.
- Alperin, Juan Pablo, Jason Portenoy, Kyle Demes, Vincent Larivière, and Stefanie Haustein (2024). "An analysis of the suitability of OpenAlex for bibliometric analyses". In: *arXiv preprint arXiv:2404.17663*.
- Babcock, Linda, Maria P. Recalde, Lise Vesterlund, and Laurie Weingart (Mar. 2017). "Gender Differences in Accepting and Receiving Requests for Tasks with Low Promotability". In: *American Economic Review* 107.3, pp. 714–47. DOI: 10.1257/aer.20141734. URL: <https://www.aeaweb.org/articles?id=10.1257/aer.20141734>.
- Bagues, Manuel, Milan Makany, Giulia Vattuone, and Natalia Zinovyeva (2026). "Female Promotions and the Academic Pipeline: Evidence from a Natural Experiment". In: *IZA Discussion Paper*.
- Bagues, Manuel and Maria J Perez-Villadoniga (2013). "Why do I like people like me?" In: *Journal of Economic Theory* 148.3, pp. 1292–1299.
- Bagues, Manuel, Mauro Sylos-Labini, and Natalia Zinovyeva (Apr. 2017). "Does the Gender Composition of Scientific Committees Matter?" In: *American Economic Review* 107.4, pp. 1207–38. DOI: 10.1257/aer.20151211. URL: <https://www.aeaweb.org/articles?id=10.1257/aer.20151211>.
- Baltrunaite, Audinga, Piera Bello, Alessandra Casarico, and Paola Profeta (2014). "Gender quotas and the quality of politicians". In: *Journal of Public Economics* 118, pp. 62–74.
- Baron, Justus, Bernhard Ganglmair, Nicola Persico, Timothy Simcoe, and Emanuele Tarantino (2024). "Representation is not sufficient for selecting gender diversity". In: *Research Policy* 53.6, p. 104994.
- Bello, Piera, Alessandra Casarico, and Debora Nozza (2025). "Research Similarity and Women in Academia". In: *The Economic Journal*, ueaf113.
- Bertrand, Marianne, Sandra E Black, Sissel Jensen, and Adriana Lleras-Muney (June 2014). *Breaking the Glass Ceiling? The Effect of Board Quotas on Female Labor Market Outcomes in Norway*. Working Paper 20256. National Bureau of Economic Research. DOI: 10.3386/w20256. URL: <http://www.nber.org/papers/w20256>.
- Blau, Francine D and Lawrence M Kahn (2017). "The gender wage gap: Extent, trends, and explanations". In: *Journal of economic literature* 55.3, pp. 789–865.

Born, Andreas, Eva Ranehill, and Anna Sandberg (2022). "Gender and willingness to lead: Does the gender composition of teams matter?" In: *Review of Economics and Statistics* 104.2, pp. 259–275.

Boschini, Anne and Anna Sjögren (2007). "Is team formation gender neutral? Evidence from coauthorship patterns". In: *Journal of Labor Economics* 25.2, pp. 325–365.

Bosquet, Clément, Pierre-Philippe Combes, and Cecilia García-Peñalosa (2019). "Gender and promotions: Evidence from academic economists in France". In: *The Scandinavian Journal of Economics* 121.3, pp. 1020–1053.

Carlsson, Magnus, Henning Finseraas, Arnfinn H Midtbøen, and Guðbjörg Linda Rafnsdóttir (2021). "Gender bias in academic recruitment? Evidence from a survey experiment in the Nordic region". In: *European Sociological Review* 37.3, pp. 399–410.

Ceci, Stephen J., Donna K. Ginther, Shulamit Kahn, and Wendy M. Williams (2014). "Women in Academic Science: A Changing Landscape". In: *Psychological Science in the Public Interest* 15.3. PMID: 26172066, pp. 75–141. DOI: 10.1177/1529100614541236. eprint: <https://doi.org/10.1177/1529100614541236>. URL: <https://doi.org/10.1177/1529100614541236>.

Ceci, Stephen J. and Wendy M. Williams (2011). "Understanding current causes of women's underrepresentation in science". In: *Proceedings of the National Academy of Sciences* 108.8, pp. 3157–3162. ISSN: 0027-8424. DOI: 10.1073/pnas.1014871108. eprint: <https://www.pnas.org/content/108/8/3157.full.pdf>. URL: <https://www.pnas.org/content/108/8/3157>.

Cecchi, Daniele, Simona Cicognani, and Nevena Kulic (2019). "Gender quotas or girls' networks? Evidence from an Italian research selection". In: *Work, Employment and Society* 33.3, pp. 462–482.

Coffman, Katherine B, Christine L Exley, and Muriel Niederle (2021). "The role of beliefs in driving gender discrimination". In: *Management Science* 67.6, pp. 3551–3569.

Cosentino, Sarah (2025). "Gender inequality in Japanese academia". In: *Nature Reviews Electrical Engineering*, pp. 1–2.

De Acutis, Costanza, Andrea Weber, and Elisabeth Wurm (2024). "The effects of board gender quotas: A meta-analysis". In: *Labour Economics* 91, p. 102634.

De Paola, Maria, Michela Ponzo, and Vincenzo Scoppa (2018). "Are men given priority for top jobs? Investigating the glass ceiling in Italian academia". In: *Journal of Human Capital* 12.3, pp. 475–503.

De Paola, Maria and Vincenzo Scoppa (2015). "Gender Discrimination and Evaluators' Gender: Evidence from Italian Academia". In: *Economica* 82.325, pp. 162–188. DOI: <https://doi.org/10.1111/ecca.12107>. eprint: <https://onlinelibrary>.

REFERENCES

167

wiley.com/doi/pdf/10.1111/ecca.12107. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/ecca.12107>.

Deschamps, Pierre (2024). "Gender quotas in hiring committees: A boon or a bane for women?" In: *Management Science* 70.11, pp. 7486–7505.

Dolado, Juan, Florentino Felgueroso, and Miguel Almunia (Sept. 2012). "Are men and women-economists evenly distributed across research fields? Some new empirical evidence". In: *SERIEs: Journal of the Spanish Economic Association* 3.3, pp. 367–393. DOI: 10.1007/s13209-011-0065-4. URL: <https://ideas.repec.org/a/spr/series/v3y2012i3p367-393.html>.

Esnard, Catherine and Roberta de R. Grangeiro (2025). "Cultural Barriers to Women's Progression in Academic Careers: A France-Brazil Comparison Through the Lens of the Queen Bee Phenomena". In: *Scandinavian Journal of Psychology* 66, pp. 210–218. DOI: 10.1111/sjop.13078.

European Commission (2018). *She Figures 2018*. Brussels: European Commission. — (2024). *She Figures 2024*. Insights on gender equality in research and innovation. Brussels: European Commission.

European Parliament and Council of the European Union (2022). *Directive (EU) 2022/2381 of the European Parliament and of the Council of 23 November 2022 on Improving the Gender Balance among Directors of Listed Companies and Related Measures*. Official Journal of the European Union, L 315. 7 December 2022, pp. 44–59. URL: <https://eur-lex.europa.eu/eli/dir/2022/2381/oj>.

Exley, Christine L, Muriel Niederle, and Lise Vesterlund (2020). "Knowing when to ask: The cost of leaning in". In: *Journal of Political Economy* 128.3, pp. 816–854.

Gërxxhani, Klarita, Nevena Kulic, Alessandra Rusconi, and Heike Solga (2025). "Gender bias in evaluating assistant professorship applicants? Evidence from harmonized survey experiments in Germany and Italy". In: *Social Science Research* 126, p. 103113.

Ginther, Donna K. and Shulamit Kahn (Nov. 2006). *Does Science Promote Women? Evidence from Academia 1973-2001*. Working Paper 12691. National Bureau of Economic Research.

Goldin, Claudia (2014). "A grand gender convergence: Its last chapter". In: *American economic review* 104.4, pp. 1091–1119.

Harris, Richard, Mariluz Mate-Sanchez-Val, and Manuel Ruiz Marín (2025). "Gender disparities in promotions and exiting in UK Russell Group universities". In: *Applied Economics* 57.31, pp. 4441–4457.

Hilmer, Christiana and Michael Hilmer (2007). "Women helping women, men helping women? Same-gender mentoring, initial job placements, and early career publishing success for economics PhDs". In: *American Economic Review* 97.2, pp. 422–426.

Huang, Junming, Alexander J Gates, Roberta Sinatra, and Albert-László Barabási (2020). "Historical comparison of gender inequality in scientific careers across countries and disciplines". In: *Proceedings of the national academy of sciences* 117.9, pp. 4609–4616.

Iaria, Alessandro, Carlo Schwarz, and Fabian Waldinger (2022). *Gender gaps in academia: Global evidence over the twentieth century*. Centre for Economic Policy Research.

Kam, Jihye and Soohyung Lee (2024). "Female representation in academia: are women underrepresented among faculty at public colleges and universities?" In: *International Economic Journal* 38.2, pp. 185–211.

Karpowitz, Christopher F, Stephen D O'Connell, Jessica Preece, and Olga Stoddard (2024). "Strength in numbers? gender composition, leadership, and women's influence in teams". In: *Journal of Political Economy* 132.9, pp. 3077–3114.

Kim, Daehyun and Laura T Starks (2016). "Gender diversity on corporate boards: Do women contribute unique skills?" In: *American Economic Review* 106.5, pp. 267–271.

Marini, Giulio and Viviana Meschitti (2018). "The trench warfare of gender discrimination: evidence from academic promotions to full professor in Italy". In: *Scientometrics* 115.2, pp. 989–1006.

Morettini, Lucio and Massimiliano Tani (2025). "Gender and career progression in academia: European evidence". In: *European Journal of Higher Education*, pp. 1–23.

Morgan, Allison C, Samuel F Way, Michael JD Hoefer, Daniel B Larremore, Mirta Galesic, and Aaron Clauset (2021). "The unequal impact of parenthood in academia". In: *Science advances* 7.9, eabd1996.

Moss-Racusin, Corinne A, John F Dovidio, Victoria L Brescoll, Mark J Graham, and Jo Handelsman (2012). "Science faculty's subtle gender biases favor male students". In: *Proceedings of the national academy of sciences* 109.41, pp. 16474–16479.

Müller, Jörg and Rachel Palmén (2025). "Pathways to gender equality in academia. National factors shaping women's early career progression across Europe". In: *Science and Public Policy*.

Nakajima, Kazuki, Ruodan Liu, Kazuyuki Shudo, and Naoki Masuda (2023). "Quantifying gender imbalance in East Asian academia: Research career and citation practice". In: *Journal of informetrics* 17.4, p. 101460.

OECD (2025). *Gender Equality in a Changing World: Taking Stock and Moving Forward*. Gender Equality at Work. Paris: OECD Publishing. DOI: [10.1787/e808086f-en](https://doi.org/10.1787/e808086f-en). URL: <https://doi.org/10.1787/e808086f-en>.

REFERENCES

169

- Osterloh, Margit and Katja Rost (2025). "Exploring the leaky pipeline: Tokenism, status group effects, or self-selection?" In: *European Management Review*.
- Priem, Jason, Heather Piwowar, and Richard Orr (2022). "OpenAlex: A fully-open index of scholarly works, authors, venues, institutions, and concepts". In: *arXiv preprint arXiv:2205.01833*.
- Radbruch, Jonas and Amelie Schiprowski (2023). *Committee deliberation and gender differences in influences*. Tech. rep. Discussion Paper.
- Spoon, Katie, Nicholas LaBerge, K Hunter Wapman, Sam Zhang, Allison C Morgan, Mirta Galesic, Bailey K Fosdick, Daniel B Larremore, and Aaron Clauset (2023). "Gender and retention patterns among US faculty". In: *Science advances* 9.42, eadi2205.
- Steinpreis, Rhea E, Katie A Anders, and Dawn Ritzke (1999). "The impact of gender on the review of the curricula vitae of job applicants and tenure candidates: A national empirical study". In: *Sex roles* 41.7, pp. 509–528.
- Williams, Wendy M. and Stephen J. Ceci (2015). "National hiring experiments reveal 2:1 faculty preference for women on STEM tenure track". In: *Proceedings of the National Academy of Sciences* 112.17, pp. 5360–5365. ISSN: 0027-8424. DOI: 10.1073/pnas.1418878112. eprint: <https://www.pnas.org/content/112/17/5360.full.pdf>. URL: <https://www.pnas.org/content/112/17/5360>.
- Zinovyeva, Natalia and Manuel Bagues (Apr. 2015). "The Role of Connections in Academic Promotions". In: *American Economic Journal: Applied Economics* 7.2, pp. 264–92. DOI: 10.1257/app.20120337. URL: <https://www.aeaweb.org/articles?id=10.1257/app.20120337>.

Appendix

Appendix 3.A Regulatory History and National Rules

3.A.1 Regulatory Changes and Ongoing Reform

The Italian academic recruitment system has undergone and is still undergoing profound transformations over the last twenty years, and it is therefore useful to provide a historical context for the regulations in place during the reform period analyzed in this chapter. Similar analyses on the effect of committees have been studied, for example, by De Paola and Scoppa, 2015¹⁶, but under the regulations preceding those analyzed here.

Until 2008, competitions for associate and full professor positions were organized at the local level by each university, but the four external committee members were elected by professors in the field on a national basis, a mechanism widely criticized for its permeability to clientelistic dynamics and agreements among academic barons. In 2008, during a final round of competitions held under the old regime, the government intervened by establishing that the four external members should be drawn by lot by the Ministry from among all full professors in the field, introducing for the first time an element of exogenous randomization in the composition of committees. Law 240/2010 then redesigned the system, shifting competitions entirely to the local level and introducing the National Scientific Habilitation (ASN) as a prerequisite for accessing local selection procedures. Under this new regime, which constitutes the institutional framework of the present chapter, each university autonomously organizes its own selection procedures and appoints committees according to internal rules. It is in this context that the University of Pisa introduced in July 2019 the reform under analysis, imposing a minimum quota of one third for each gender and an internal lottery mechanism to make this quota binding.

The system is currently undergoing further transition. Legislative Bill 2735, approved by the Senate in December 2025 and under examination by the Chamber of Deputies at the time of writing of this chapter¹⁷, provides for the replacement of the ASN through a system of self-certification of minimum requirements and rede-

¹⁶This is the institutional context analyzed by De Paola and Scoppa, 2015, who exploit this random assignment to causally identify the effect of the presence of female evaluators on the probability of success of female candidates, finding that all-male committees reduce women's chances by approximately 6–7 percentage points, an effect that disappears with the introduction of at least one female member.

¹⁷Legislative Bill A.C. 2735, *Revisione delle modalità di accesso, valutazione e reclutamento del personale ricercatore e docente universitario*, XIX Legislatura. Transmitted by the Senate on 9 December 2025, currently under examination by the VII Culture Committee. Available at <https://camera.it/leg19/126?idDocumento=2735>.

3.A. REGULATORY HISTORY AND NATIONAL RULES

171

finest the composition of committees: one internal member appointed by the university and four external members drawn by lot from national lists of candidates prepared by the Ministry, in compliance with the principle of gender balance. If approved, this bill would represent a partial return to the centralized mechanism, albeit with important differences, in particular the lottery from lists of pre-selected volunteers rather than from the universe of professors in the field.

3.A.2 A National Survey on Committee Composition Rules

To place the Pisa reform in context, we assembled information on how other Italian universities regulate the gender composition of selection committees, through two complementary exercises.

First, we surveyed the Equal Opportunity Committees (CUG) of Italian universities in 2023. Forty-one universities responded with usable information. Among them, 31 report some form of gender quota in the composition of selection committees, while 10 provide for no such requirement. Random selection of committee members is used by 25 of the 41 respondents. Among the universities that specify a numerical threshold, the quota ranges from a minimum of 20 percent (Basilicata) to a maximum of 50 percent (Chieti-Pescara), with the large majority set around one third. Only seven universities combine a one-third quota with a random-selection mechanism, which makes them the closest institutional analogues to the Pisa reform: besides Pisa itself, these are Bologna, Bari Aldo Moro, Bergamo, the University for Foreigners of Perugia, the Scuola Normale Superiore, and the IMT School for Advanced Studies Lucca.

Second, we examined the internal regulations governing selection procedures at a broader set of universities, to check whether a gender-balance rule appears in the formal rules even where it was not reported in the survey. Here the regulations are, in most cases, not precise enough to classify. They typically refer to gender balance in aspirational terms, with formulations such as “where possible” or “preferably, both genders should be represented,” without specifying either a binding quota or a random-selection mechanism. We therefore do not attempt a systematic count from the regulations and rely on the survey responses, which distinguish quotas from random selection explicitly.

Taken together, these results show that attention to gender balance in selection committees remains largely at the discretion of individual universities, both in whether it is written into the rules at all and, still more, in whether it is translated into a binding mechanism. The combination Pisa adopts, a one-third quota enforced through a lottery, is rare.

Table Appendix

3.A.3 Summary Statistics and descriptive

Table A1: Competitions, candidates and gender shares by year

Year	N. Calls	Share F. Comm.	N. Candidates	Share F. Cand.	Share F. Winners
2015	30	0.14	141	0.25	0.27
2016	97	0.18	255	0.26	0.30
2018	13	0.03	96	0.14	0.00
2019 (pre)	67	0.37	292	0.33	0.43
2019 (post)	67	0.40	310	0.28	0.36
2020	57	0.43	231	0.33	0.32
2021	165	0.42	689	0.28	0.34
2022	133	0.41	505	0.34	0.39
2023	100	0.45	382	0.30	0.32
2024	60	0.48	249	0.38	0.50
Total	789	0.37	3150	0.30	0.35

Table A2: Competitions, candidates and gender shares by position type

Position type	N. Calls	Share F. Comm.	N. Candidates	Share F. Cand.	Share F. Winners
Professore I fascia	253	0.34	885	0.24	0.30
Professore II fascia	154	0.29	730	0.26	0.37
Ricercatori	382	0.43	1535	0.35	0.38
Total	789	0.37	3150	0.30	0.35

3.A. REGULATORY HISTORY AND NATIONAL RULES

173

Table A3: Descriptive statistics: candidate, committee, and competition characteristics


Variable	Overall	Female	Male
Winner	0.25	0.29	0.23
Pisa affiliated	0.28	0.32	0.26
N. candidates	7.07 (6.03)	5.74 (4.34)	7.65 (6.55)
Share female candidates	0.30 (0.29)	0.57 (0.28)	0.18 (0.19)
Share female comm. (actual)	0.38 (0.25)	0.43 (0.25)	0.36 (0.25)
Share female comm. (expected)	0.38 (0.20)	0.43 (0.19)	0.36 (0.20)
Documents (at comp.)	108.4 (203.1) [51.0]	83.9 (162.9) [39.5]	118.9 (217.4) [56.0]
Citations (at comp.)	4259 (14482) [749]	3231 (12016) [473]	4703 (15406) [942]
Documents (current)	167.5 (277.2) [86.0]	128.3 (203.8) [65.0]	184.2 (301.5) [94.0]
Citations (current)	5514 (16649) [1364]	4139 (13725) [1007]	6098 (17717) [1497]
H-index (current)	23.9 (24.0) [19.0]	20.7 (20.8) [16.0]	25.3 (25.1) [20.0]
Avg. H-index (comm.)	29.6 (20.0) [28.4]	28.8 (20.2) [28.3]	30.0 (20.0) [28.4]
Avg. citations (comm.)	5701 (8420) [3235]	5367 (7351) [3217]	5845 (8839) [3245]
Avg. documents (comm.)	193.3 (173.5) [159.0]	178.8 (152.5) [151.8]	199.5 (181.5) [162.6]
Avg. age (comm.)	57.7 (4.6)	58.1 (4.8)	57.6 (4.5)
% missing: documents	0.13	0.14	0.12
% missing: citations	0.13	0.14	0.12
N	3150	950	2200

Notes: Cell entries report the mean, with the standard deviation in parentheses below and, where relevant, the median in square brackets. "At comp." denotes the value up to the competition year, "current" the value at the time of data collection.

Table A4: Share of female evaluators by CUN area and year

CUN Area	2015	2016	2018	2019 _{pre}	2019 _{post}	2020	2021	2022	2023	2024
Agricultural Sciences	0.00	0.08		0.44	0.27	0.40	0.44	0.43	0.50	0.56
Architecture		0.48	0.00	0.50	0.47	0.56	0.43	0.59	0.58	0.50
Biology	0.33	0.33		0.41	0.27	0.53	0.43	0.44	0.76	0.60
Chemistry		0.40		0.33	0.53	0.33	0.49	0.57	0.53	0.33
Civil Engineering		0.00		0.33	0.60	0.40	0.67	0.57	0.40	
Earth Sciences	0.00			0.33	0.33	0.22	0.22	0.17		0.17
Economics	0.33	0.25		0.38	0.42	0.50	0.48	0.34	0.53	0.44
History		0.00	0.00	0.42	0.27	0.67	0.51	0.50	0.38	0.42
Industrial Engineering	0.07	0.17	0.00	0.20	0.31	0.40	0.41	0.33	0.37	0.40
Law	0.00	0.22		0.60	0.40	0.22	0.42	0.41	0.30	0.67
Mathematics and Computer Science	0.00	0.03	0.17	0.33	0.54	0.40	0.29	0.35	0.53	0.40
Medicine	0.20	0.02	0.00	0.48	0.32	0.48	0.37	0.37	0.42	0.48
Physics	0.12	0.05		0.00	0.53	0.56	0.38	0.38	0.29	0.83
Political Science	0.33				0.33		0.52	0.33	0.67	

Figure A1: Average share of female committee members by year.



Notes: Average realized female committee share by year; 2019 split into pre- and post-reform. The dashed line marks one third. 2017 is excluded.

3.A. REGULATORY HISTORY AND NATIONAL RULES

175

3.A.4 Post Reform Analysis

Table A5: Committee gender composition and the probability of winning, excluding researcher competitions with ≥ 6 candidates (candidate level, contestable sample).

	(1) No FE	(2) Year	(3) CUN	(4) SSD	(5) Yr+CUN	(6) Yr+SSD
Female $\times$ Share female (actual)	−0.071 (0.120)	−0.103 (0.123)	−0.062 (0.122)	−0.042 (0.145)	−0.088 (0.125)	−0.068 (0.143)
Num. Obs.	1423	1423	1420	1423	1420	1423

Notes: LPM at the candidate level; the outcome equals one if the candidate wins. The coefficient of interest is the interaction of female candidacy with the actual female committee share. Sample: contestable competitions ($n > 1$), excluding researcher competitions with six or more observed candidates. Controls as in Table 3.7 (not shown). SE clustered at the SSD level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A6: Main specification with all coefficients (contestable sample, $n > 1$).

	(1) No FE	(2) Year	(3) CUN	(4) SSD	(5) Year + CUN	(6) Year + SSD
Female candidate	0.007 (0.052)	0.008 (0.052)	-0.004 (0.052)	-0.008 (0.061)	-0.001 (0.053)	-0.009 (0.062)
Share female (actual)	-0.043 (0.029)	-0.024 (0.031)	-0.052* (0.027)	-0.094*** (0.031)	-0.031 (0.029)	-0.067* (0.036)
Female × Share female (actual)	0.038 (0.072)	0.019 (0.075)	0.046 (0.073)	0.051 (0.080)	0.028 (0.076)	0.037 (0.082)
Share female (expected)	0.057 (0.046)	0.028 (0.046)	0.055 (0.047)	0.083* (0.047)	0.026 (0.048)	0.065 (0.053)
Female × Share female (expected)	-0.006 (0.120)	0.010 (0.121)	-0.007 (0.123)	0.010 (0.141)	0.004 (0.124)	0.025 (0.144)
Number of candidates	-0.009** (0.004)	-0.010*** (0.003)	-0.009** (0.004)	-0.008** (0.004)	-0.009** (0.003)	-0.008** (0.003)
Pisa candidate	0.281*** (0.032)	0.286*** (0.032)	0.284*** (0.033)	0.324*** (0.038)	0.289*** (0.033)	0.327*** (0.038)
Committee avg. age	0.000 (0.001)	-0.001 (0.001)	-0.002* (0.001)	-0.002 (0.002)	-0.003* (0.001)	-0.002 (0.002)
Committee avg. log cites	0.001 (0.006)	0.003 (0.006)	0.001 (0.007)	0.013 (0.010)	0.003 (0.007)	0.015 (0.011)
Committee avg. log docs	-0.017 (0.015)	-0.019 (0.016)	-0.016 (0.018)	-0.039 (0.032)	-0.016 (0.018)	-0.036 (0.034)
Log docs (until competition)	0.040*** (0.013)	0.041*** (0.014)	0.038** (0.015)	0.052*** (0.016)	0.038** (0.015)	0.052*** (0.017)
Log cites (until competition)	-0.017** (0.007)	-0.017** (0.007)	-0.016* (0.008)	-0.019** (0.009)	-0.015* (0.008)	-0.019** (0.009)
Share Pisa committee (actual)	0.010 (0.055)	-0.016 (0.046)	-0.008 (0.064)	0.473 (0.683)	-0.032 (0.052)	0.187 (0.631)
Share Pisa committee (expected)	0.003 (0.129)	0.024 (0.124)	0.069 (0.132)	-0.725 (1.155)	0.094 (0.120)	-0.269 (1.078)
Committee docs missing	-0.045 (0.054)	-0.046 (0.056)	-0.059 (0.082)		-0.041 (0.080)	
Candidate cites missing	0.108*** (0.037)	0.108*** (0.038)	0.103** (0.040)	0.156*** (0.047)	0.103** (0.041)	0.159*** (0.047)
Num.Obs.	2240	2240	2240	2240	2240	2240
R2	0.132	0.135	0.136	0.173	0.139	0.175

Notes: Same specification as the main table, with all coefficients displayed: candidate and committee controls and missing-value indicators. LPM with progressive fixed effects. SE clustered at the SSD level in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.A. REGULATORY HISTORY AND NATIONAL RULES

177

Table A7: LPM with progressive fixed effects (full sample).

	(1) No FE	(2) Year	(3) CUN	(4) SSD	(5) Year + CUN	(6) Year + SSD
Female candidate	-0.022 (0.054)	-0.020 (0.055)	-0.036 (0.054)	-0.017 (0.060)	-0.033 (0.055)	-0.017 (0.061)
Share female (actual)	-0.035 (0.034)	-0.013 (0.036)	-0.045 (0.032)	-0.089** (0.034)	-0.021 (0.033)	-0.061 (0.039)
Female × Share female (actual)	0.019 (0.076)	-0.005 (0.080)	0.032 (0.077)	0.024 (0.079)	0.009 (0.079)	0.006 (0.081)
Share female (expected)	0.012 (0.059)	-0.031 (0.059)	0.020 (0.056)	0.129** (0.056)	-0.022 (0.058)	0.106* (0.059)
Female × Share female (expected)	0.058 (0.127)	0.081 (0.128)	0.056 (0.130)	0.050 (0.142)	0.073 (0.130)	0.067 (0.144)
Number of candidates	-0.015** (0.006)	-0.015*** (0.005)	-0.013** (0.006)	-0.012** (0.005)	-0.013** (0.005)	-0.012** (0.005)
Log docs (until competition)	0.049*** (0.015)	0.052*** (0.015)	0.047*** (0.015)	0.059*** (0.017)	0.048*** (0.015)	0.059*** (0.017)
Log cites (until competition)	-0.020*** (0.008)	-0.021*** (0.008)	-0.020** (0.009)	-0.022** (0.010)	-0.019** (0.009)	-0.021** (0.010)
Num.Obs.	2366	2366	2366	2333	2366	2333
R2	0.176	0.182	0.185	0.237	0.189	0.240

Notes: LPM, full sample. SE clustered at the SSD level in parentheses. Controls as in Table 3.7.
 * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A8: LPM with progressive fixed effects (both-genders sample).

	(1) No FE	(2) Year	(3) CUN	(4) SSD	(5) Year + CUN	(6) Year + SSD
Female candidate	0.014 (0.062)	0.013 (0.062)	0.006 (0.062)	-0.002 (0.066)	0.007 (0.063)	-0.003 (0.067)
Share female (actual)	-0.041 (0.040)	-0.029 (0.040)	-0.051 (0.036)	-0.117*** (0.040)	-0.034 (0.037)	-0.083** (0.040)
Female × Share female (actual)	0.055 (0.084)	0.042 (0.086)	0.064 (0.084)	0.065 (0.091)	0.051 (0.086)	0.052 (0.093)
Share female (expected)	0.075 (0.053)	0.049 (0.055)	0.082 (0.054)	0.051 (0.055)	0.054 (0.056)	0.042 (0.062)
Female × Share female (expected)	-0.015 (0.142)	-0.003 (0.144)	-0.016 (0.145)	-0.009 (0.152)	-0.008 (0.147)	0.005 (0.155)
Number of candidates	-0.007** (0.003)	-0.007*** (0.003)	-0.006** (0.003)	-0.008** (0.003)	-0.006*** (0.002)	-0.009*** (0.003)
Log docs (until competition)	0.034** (0.014)	0.036** (0.015)	0.032** (0.016)	0.039** (0.018)	0.033** (0.016)	0.041** (0.018)
Log cites (until competition)	-0.017** (0.008)	-0.018** (0.008)	-0.016* (0.009)	-0.017* (0.010)	-0.016* (0.009)	-0.017* (0.010)
Num.Obs.	1688	1688	1688	1688	1688	1688
R2	0.126	0.129	0.130	0.166	0.133	0.168

Notes: LPM. Sample restricted to competitions with at least one man and one woman among candidates. SE clustered at the SSD level in parentheses. Controls as in Table 3.7. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A9: Logit with fixed effects (contestable sample).

	(1) No FE	(2) Year	(4) SSD	(6) Year + SSD
Female candidate	−0.054 (0.343)	−0.045 (0.347)	−0.049 (0.456)	−0.053 (0.463)
Share female (actual)	−0.191 (0.201)	−0.170 (0.204)	−0.458* (0.261)	−0.390 (0.266)
Female × Share female (actual)	0.222 (0.497)	0.188 (0.508)	0.282 (0.611)	0.251 (0.623)
Share female (expected)	0.114 (0.344)	0.032 (0.345)	0.379 (0.441)	0.298 (0.448)
Female × Share female (expected)	0.133 (0.808)	0.155 (0.820)	0.157 (1.059)	0.187 (1.079)
Number of candidates	−0.170*** (0.029)	−0.166*** (0.029)	−0.131*** (0.038)	−0.126*** (0.037)
Log docs (until competition)	0.311*** (0.097)	0.319*** (0.100)	0.426*** (0.135)	0.429*** (0.138)
Log cites (until competition)	−0.131*** (0.050)	−0.132*** (0.051)	−0.153** (0.073)	−0.149** (0.074)
Num.Obs.	2240	2240	2240	2240

Notes: Logit estimated with `feglm`; coefficients on the log-odds scale. SE clustered at the SSD level in parentheses. Controls as in Table 3.7. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.



3.A. REGULATORY HISTORY AND NATIONAL RULES

179

Table A10: OLS progressive specification (contestable sample, $n > 1$).

	(1)	(2)	(3)	(4)	(5)	(6)
Female candidate	0.056*** (0.019)	0.029 (0.019)	0.030 (0.019)	0.007 (0.041)	0.006 (0.038)	0.007 (0.055)
Share female (actual)			-0.027 (0.035)	-0.042 (0.041)	-0.024 (0.040)	-0.043 (0.047)
Female × Share female (actual)				0.050 (0.079)	0.037 (0.074)	0.038 (0.093)
Share female (expected)						0.057 (0.075)
Female × Share female (expected)						-0.006 (0.145)
Number of candidates		-0.012*** (0.001)	-0.012*** (0.001)	-0.012*** (0.001)	-0.009*** (0.001)	-0.009*** (0.001)
Log docs (until competition)					0.040*** (0.015)	0.040*** (0.015)
Log cites (until competition)					-0.017** (0.008)	-0.017** (0.008)
Num.Obs.	2240	2240	2240	2240	2240	2240
R2	0.004	0.041	0.041	0.041	0.132	0.132

Notes: OLS, controls added progressively across columns. Heteroskedasticity-robust (HC1) SE in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A11: OLS progressive specification (full sample).

	(1)	(2)	(3)	(4)	(5)	(6)
Female candidate	0.059*** (0.020)	0.020 (0.019)	0.022 (0.019)	-0.005 (0.041)	-0.005 (0.038)	-0.022 (0.054)
Share female (actual)			-0.045 (0.036)	-0.063 (0.042)	-0.031 (0.041)	-0.035 (0.048)
Female × Share female (actual)				0.059 (0.080)	0.042 (0.074)	0.019 (0.094)
Share female (expected)						0.012 (0.077)
Female × Share female (expected)						0.058 (0.146)
Number of candidates		-0.018*** (0.001)	-0.018*** (0.001)	-0.018*** (0.001)	-0.015*** (0.001)	-0.015*** (0.001)
Log docs (until competition)					0.049*** (0.015)	0.049*** (0.015)
Log cites (until competition)					-0.020** (0.008)	-0.020** (0.008)
Num.Obs.	2366	2366	2366	2366	2366	2366
R2	0.004	0.074	0.075	0.075	0.176	0.176

Notes: OLS, full sample, controls added progressively across columns. Heteroskedasticity-robust (HC1) SE in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A12: LPM with progressive fixed effects, rank controls (contestable sample).

	(1) No FE	(2) Year	(3) CUN	(4) SSD	(5) Year + CUN	(6) Year + SSD
Female candidate	0.008 (0.052)	0.010 (0.052)	-0.004 (0.053)	-0.008 (0.061)	-0.001 (0.053)	-0.009 (0.062)
Share female (actual)	-0.044 (0.029)	-0.024 (0.031)	-0.051* (0.027)	-0.084*** (0.031)	-0.030 (0.030)	-0.060 (0.037)
Female × Share female (actual)	0.036 (0.072)	0.018 (0.075)	0.046 (0.073)	0.050 (0.081)	0.028 (0.076)	0.038 (0.083)
Share female (expected)	0.055 (0.046)	0.028 (0.046)	0.054 (0.048)	0.066 (0.049)	0.025 (0.049)	0.051 (0.054)
Female × Share female (expected)	-0.013 (0.120)	0.002 (0.121)	-0.009 (0.123)	0.007 (0.142)	0.000 (0.124)	0.018 (0.144)
Number of candidates	-0.009** (0.004)	-0.009*** (0.003)	-0.009** (0.004)	-0.007* (0.004)	-0.009** (0.003)	-0.008** (0.004)
Num.Obs.	2240	2240	2240	2240	2240	2240
R2	0.135	0.138	0.140	0.176	0.143	0.178

Notes: LPM. Candidate productivity entered as within-competition percentile ranks rather than logs. SE clustered at the SSD level in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A13: LPM with competition-type fixed effects (contestable sample).

	(7) CompTitle	(8) Year+CompTitle	(9) Yr+CUN+CompTitle
Female candidate	0.005 (0.051)	0.008 (0.052)	0.000 (0.052)
Share female (actual)	-0.015 (0.031)	0.001 (0.034)	-0.010 (0.032)
Female × Share female (actual)	0.020 (0.074)	0.004 (0.077)	0.015 (0.077)
Share female (expected)	0.019 (0.048)	0.001 (0.049)	0.003 (0.050)
Female × Share female (expected)	0.015 (0.121)	0.025 (0.122)	0.018 (0.125)
Number of candidates	-0.010** (0.004)	-0.010*** (0.004)	-0.010*** (0.004)
Log docs (until competition)	0.046*** (0.014)	0.045*** (0.014)	0.040*** (0.016)
Log cites (until competition)	-0.018** (0.007)	-0.018** (0.007)	-0.015* (0.008)
Num.Obs.	2240	2240	2240
R2	0.135	0.139	0.142

Notes: LPM with competition-type fixed effects (PO / PA / RIC / RTT). SE clustered at the SSD level in parentheses. Controls as in Table 3.7. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.A. REGULATORY HISTORY AND NATIONAL RULES

181

Table A14: Clustering robustness (preferred FE, contestable sample).

	SSD cluster	Competition cluster
Female candidate	−0.009 (0.062)	−0.009 (0.063)
Share female (actual)	−0.067* (0.036)	−0.067* (0.039)
Female × Share female (actual)	0.037 (0.082)	0.037 (0.096)
Share female (expected)	0.065 (0.053)	0.065 (0.067)
Female × Share female (expected)	0.025 (0.144)	0.025 (0.167)
Number of candidates	−0.008** (0.003)	−0.008*** (0.002)
Log docs (until competition)	0.052*** (0.017)	0.052*** (0.019)
Log cites (until competition)	−0.019** (0.009)	−0.019* (0.010)
Num.Obs.	2240	2240
R2	0.175	0.175

Notes: LPM with year and SSD fixed effects. Columns differ only in the level at which SE are clustered. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A15: Heterogeneity by field: STEM vs non-STEM (preferred FE, contestable sample).

	STEM	Non-STEM
Female candidate	0.016 (0.072)	−0.057 (0.103)
Share female (actual)	−0.018 (0.038)	−0.092 (0.091)
Female × Share female (actual)	−0.055 (0.091)	0.118 (0.154)
Share female (expected)	0.023 (0.054)	0.037 (0.125)
Female × Share female (expected)	0.147 (0.170)	−0.045 (0.243)
Number of candidates	−0.007** (0.003)	−0.026*** (0.007)
Log docs (until competition)	0.048** (0.018)	0.060* (0.032)
Log cites (until competition)	−0.015 (0.012)	−0.026* (0.015)
Num.Obs.	1367	873
R2	0.173	0.185

Notes: LPM with year and SSD fixed effects. SE clustered at the SSD level in parentheses. STEM corresponds to CUN areas 1–5, 8, 9. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.A. REGULATORY HISTORY AND NATIONAL RULES

183

Table A16: Coefficient on $Female \times Share\ female\ (actual)$ and minimum detectable effects

Sample	Fixed effects	$\hat{\beta}$	Clustered SE	MDE (p.p.)
Full	No FE	0.019	0.076	21.3
	Year	−0.005	0.080	22.3
	SSD	0.024	0.079	22.1
	Year + SSD	0.006	0.081	22.6
Contestable ($n > 1$)	No FE	0.038	0.072	20.2
	Year	0.019	0.075	21.1
	SSD	0.051	0.080	22.4
	Year + SSD	0.037	0.082	23.0
Both genders	No FE	0.055	0.084	23.5
	Year	0.042	0.086	24.2
	SSD	0.065	0.091	25.6
	Year + SSD	0.052	0.093	26.0

Notes: Each row reports the coefficient on the interaction of interest, its standard error clustered at the SSD level, and the implied minimum detectable effect at 80% power and $\alpha = 0.05$, $MDE = 2.80 \times \widehat{SE} \times 100$.

Committee Composition and Female Winners: Competition Level

As a robustness check on independent observations, we restate the candidate-level result at the competition level. We collapse the data to one row per competition and ask whether a more female committee makes a female winner more likely, estimating a linear probability model in which the outcome equals one if the competition is won by a woman:

$$\text{FemaleWinner}_c = \gamma_0 + \gamma_1 \text{ShareFemale}_c^{\text{actual}} + \gamma_2 \text{ShareFemale}_c^{\text{expected}} + Z_c' \lambda + \mu_s + \tau_t + \nu_c, \quad (3.5)$$

where Z_c controls for the share of female candidates, an indicator for whether the candidate with the most publications is a woman, the number of candidates, and the share of Pisa-affiliated candidates.

The conclusion matches the candidate-level analysis. In the preferred specification γ_1 is 0.053 (0.097) on the contestable sample (Table A17) and −0.011 (0.078) on the full sample, in both cases indistinguishable from zero. The share of female candidates is, as expected, the dominant predictor of a female winner, with a coefficient close to one. Two caveats temper how much this design adds. First, it is less powerful than the candidate-level analysis: there are far fewer competitions than candidates, and saturating with SSD fixed effects drops singleton sectors, leaving 347 competitions in the contestable sample, so the implied minimum de-

tectable effect is around 27 percentage points. Second, the expected share enters negatively and significantly once SSD fixed effects are included, confirming that the endogenous, shortlist-driven part of committee composition is correlated with outcomes and must be conditioned out before γ_1 can be read. The only positive and marginally significant estimate appears in the both-genders sample (0.254, $p < 0.1$; Table A19), the subsample in which committee composition has the most scope to matter, but it does not survive in the other samples and we do not read it as evidence of an effect.

Table A17: Committee gender composition and the probability of a female winner (competition level, contestable sample).

	(1) No FE	(2) Year	(3) CUN	(4) SSD	(5) Year + SSD
Share female (actual)	0.079 (0.077)	0.081 (0.077)	0.063 (0.077)	0.053 (0.098)	0.053 (0.097)
Share female (expected)	-0.031 (0.135)	-0.028 (0.138)	0.038 (0.148)	-0.318* (0.172)	-0.335* (0.177)
Share female candidates	0.962*** (0.071)	0.959*** (0.071)	1.027*** (0.073)	1.013*** (0.126)	0.997*** (0.126)
Most-published candidate is female	0.091 (0.060)	0.089 (0.060)	0.092 (0.059)	0.054 (0.078)	0.057 (0.079)
Number of candidates	0.004 (0.006)	0.004 (0.006)	-0.002 (0.005)	0.004 (0.007)	0.004 (0.007)
Share Pisa candidates	0.052 (0.061)	0.048 (0.064)	0.047 (0.068)	-0.089 (0.100)	-0.091 (0.115)
Num.Obs.	456	456	456	347	347
R2	0.452	0.454	0.473	0.637	0.639

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: LPM at the competition level; the outcome equals one if the competition is won by a woman. The coefficient of interest is the actual female share of the committee. Columns differ in the set of fixed effects. SE clustered at the SSD level in parentheses. Controls included but not shown: indicator for competitions with no bibliometric data. Sample: competitions with more than one candidate. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.A. REGULATORY HISTORY AND NATIONAL RULES

185

Table A18: Committee gender composition and the probability of a female winner (competition level, full sample).

	(1) No FE	(2) Year	(3) CUN	(4) SSD	(5) Year + SSD
Share female (actual)	0.063 (0.060)	0.064 (0.060)	0.054 (0.061)	-0.006 (0.078)	-0.011 (0.078)
Share female (expected)	-0.022 (0.111)	-0.017 (0.112)	0.029 (0.119)	-0.238* (0.143)	-0.255* (0.146)
Share female candidates	0.953*** (0.055)	0.952*** (0.055)	0.993*** (0.055)	0.990*** (0.081)	0.982*** (0.082)
Most-published candidate is female	0.087 (0.054)	0.084 (0.054)	0.089* (0.054)	0.039 (0.067)	0.039 (0.068)
Number of candidates	0.004 (0.005)	0.004 (0.005)	-0.002 (0.005)	0.004 (0.006)	0.005 (0.006)
Share Pisa candidates	0.022 (0.034)	0.018 (0.036)	0.018 (0.035)	-0.027 (0.054)	-0.029 (0.060)
Num.Obs.	582	582	582	467	467
R2	0.566	0.567	0.579	0.703	0.704

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: LPM at the competition level; the outcome equals one if the competition is won by a woman. Full sample. SE clustered at the SSD level in parentheses. Controls included but not shown: indicator for competitions with no bibliometric data. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A19: Committee gender composition and the probability of a female winner (competition level, both-genders sample).

	(1) No FE	(2) Year	(3) CUN	(4) SSD	(5) Year + SSD
Share female (actual)	0.130 (0.120)	0.143 (0.120)	0.114 (0.121)	0.204 (0.148)	0.254* (0.134)
Share female (expected)	-0.084 (0.196)	-0.077 (0.198)	-0.001 (0.218)	-0.517** (0.221)	-0.484** (0.231)
Share female candidates	1.086*** (0.157)	1.082*** (0.161)	1.200*** (0.161)	1.126*** (0.276)	1.027*** (0.291)
Most-published candidate is female	0.088 (0.061)	0.085 (0.061)	0.091 (0.061)	0.068 (0.085)	0.081 (0.086)
Number of candidates	0.006 (0.007)	0.005 (0.007)	-0.003 (0.007)	0.003 (0.010)	0.003 (0.011)
Share Pisa candidates	0.095 (0.106)	0.084 (0.109)	0.086 (0.122)	-0.151 (0.200)	-0.133 (0.239)
Num.Obs.	299	299	299	204	204
R2	0.213	0.218	0.253	0.524	0.534

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: LPM at the competition level; the outcome equals one if the competition is won by a woman. Sample restricted to competitions with at least one man and one woman among candidates. SE clustered at the SSD level in parentheses. Controls included but not shown: indicator for competitions with no bibliometric data. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Robustness: excluding potentially truncated researcher competitions

Section 3.3.1 matters less here than at the candidate level. We nonetheless verify that the results are unchanged when researcher competitions with six or more candidates are excluded. The conclusion is unaffected across all three samples; in particular, the only marginally significant estimate, in the both-genders sample (Table A19), loses significance once these competitions are dropped, which reinforces our decision not to read it as evidence of an effect.

Table A20: Committee gender composition and the probability of a female winner, excluding researcher competitions with ≥ 6 candidates (competition level, contestable sample).

	(1) No FE	(2) Year	(3) CUN	(4) SSD	(5) Year + SSD
Share female (actual)	−0.070 (0.106)	−0.070 (0.105)	−0.106 (0.105)	−0.092 (0.158)	−0.097 (0.161)
Share female (expected)	0.041 (0.184)	0.044 (0.186)	0.106 (0.198)	−0.162 (0.252)	−0.173 (0.262)
Share female candidates	0.925*** (0.083)	0.918*** (0.086)	0.967*** (0.089)	1.106*** (0.118)	1.102*** (0.123)
Most-published candidate is female	0.129* (0.077)	0.133* (0.078)	0.125 (0.078)	0.045 (0.094)	0.050 (0.096)
Number of candidates	0.001 (0.005)	0.001 (0.006)	−0.008* (0.005)	0.000 (0.008)	−0.001 (0.007)
Share Pisa candidates	0.009 (0.066)	0.009 (0.068)	0.006 (0.069)	−0.109 (0.100)	−0.136 (0.105)
Num.Obs.	327	327	327	219	219
R2	0.513	0.513	0.535	0.731	0.734

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: LPM at the competition level; the outcome equals one if the competition is won by a woman. Sample: competitions with more than one candidate, excluding researcher competitions with six or more candidates. SE clustered at the SSD level in parentheses. Controls included but not shown: indicator for competitions with no bibliometric data. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.A. REGULATORY HISTORY AND NATIONAL RULES

187

Table A21: Committee gender composition and the probability of a female winner, excluding researcher competitions with ≥ 6 candidates (competition level, full sample).

	(1) No FE	(2) Year	(3) CUN	(4) SSD	(5) Year + SSD
Share female (actual)	−0.049 (0.074)	−0.050 (0.074)	−0.070 (0.073)	−0.146 (0.102)	−0.145 (0.104)
Share female (expected)	0.028 (0.135)	0.034 (0.137)	0.061 (0.144)	−0.057 (0.170)	−0.066 (0.174)
Share female candidates	0.926*** (0.062)	0.924*** (0.062)	0.943*** (0.064)	1.038*** (0.083)	1.038*** (0.086)
Most-published candidate is female	0.116* (0.066)	0.118* (0.067)	0.120* (0.067)	0.032 (0.080)	0.030 (0.081)
Number of candidates	0.001 (0.005)	0.001 (0.005)	−0.008 (0.005)	0.000 (0.007)	0.000 (0.008)
Share Pisa candidates	−0.001 (0.033)	−0.001 (0.035)	−0.011 (0.033)	−0.052 (0.049)	−0.058 (0.051)
Num.Obs.	453	453	453	333	333
R2	0.645	0.645	0.657	0.790	0.791

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: LPM at the competition level; the outcome equals one if the competition is won by a woman. Full sample, excluding researcher competitions with six or more candidates. SE clustered at the SSD level in parentheses. Controls included but not shown: indicator for competitions with no bibliometric data. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A22: Committee gender composition and the probability of a female winner, excluding researcher competitions with ≥ 6 candidates (competition level, both-genders sample).

	(1) No FE	(2) Year	(3) CUN	(4) SSD	(5) Year + SSD
Share female (actual)	−0.122 (0.195)	−0.123 (0.195)	−0.165 (0.195)	−0.013 (0.299)	0.003 (0.246)
Share female (expected)	0.042 (0.303)	0.038 (0.304)	0.093 (0.332)	−0.642 (0.384)	−0.824** (0.320)
Share female candidates	1.177*** (0.205)	1.163*** (0.213)	1.210*** (0.217)	1.158*** (0.415)	0.833* (0.425)
Most-published candidate is female	0.120 (0.080)	0.126 (0.083)	0.116 (0.083)	0.091 (0.107)	0.159 (0.108)
Number of candidates	0.005 (0.007)	0.005 (0.007)	−0.009 (0.007)	−0.010 (0.010)	−0.013 (0.011)
Share Pisa candidates	0.018 (0.125)	0.023 (0.129)	0.022 (0.143)	−0.136 (0.243)	−0.214 (0.260)
Num.Obs.	193	193	193	107	107
R2	0.231	0.233	0.279	0.658	0.676

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: LPM at the competition level; the outcome equals one if the competition is won by a woman. Sample restricted to competitions with at least one man and one woman among candidates, excluding researcher competitions with six or more candidates. SE clustered at the SSD level in parentheses. Controls included but not shown: indicator for competitions with no bibliometric data. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.A. REGULATORY HISTORY AND NATIONAL RULES

189

3.A.5 Pre-post Reform Analysis

Table A23: Treatment intensity by CUN sector: three measures of pre-reform male dominance.

CUN sector	Committee (Pisa)	National	Faculty (Pisa)
Mathematics and Computer Science	0.860	0.711	0.813
Physics	0.939	0.819	0.799
Chemistry	0.633	0.601	0.535
Earth Sciences	0.905	0.731	0.711
Biology	0.622	0.549	0.590
Medicine	0.854	0.766	0.776
Agricultural Sciences	0.792	0.680	0.628
Civil Engineering	0.833	0.732	0.852
Industrial Engineering	0.848	0.858	0.872
Architecture	0.578	0.492	0.481
History	0.815	0.574	0.676
Law	0.667	0.682	0.596
Economics	0.694	0.680	0.722
Political Science	0.667	0.650	0.771

Notes: Each measure is defined as one minus the pre-reform female share in the relevant population: members of pre-reform Pisa hiring committees (2015–2019, 2017 missing), national faculty in the field (CINECA), and faculty at the University of Pisa eligible to serve on committees. Higher values denote more male-dominated fields and hence stronger exposure to the reform. The Faculty (Pisa) measure is the preferred one used in the main analysis (Section 3.5.1).

Table A24: Pairwise correlations between the three treatment-intensity measures.

	Committee (Pisa)	National	Faculty (Pisa)
Committee (Pisa)	1.00		
National	0.78	1.00	
Faculty (Pisa)	0.75	0.82	1.00

Notes: Pairwise correlations computed across the fourteen CUN sectors reported in Table A23. The strong correlation across measures (0.75–0.82) implies that the choice among them affects precision rather than the sign of the difference-in-differences estimates.

CHAPTER 3. GENDER-BALANCED COMMITTEES

Table A25: Triple difference on winning, full sample (preferred intensity).

	No FE	CUN FE	SSD FE
Female	0.080** (0.037)	0.023 (0.035)	0.023 (0.036)
Post reform	−0.018 (0.019)	0.016 (0.020)	0.005 (0.020)
Treatment intensity	−0.039 (0.169)		
Female × Post	−0.024 (0.043)	−0.016 (0.041)	−0.007 (0.042)
Female × Intensity	0.225 (0.269)	−0.161 (0.235)	−0.175 (0.263)
Post × Intensity	−0.235 (0.172)	0.043 (0.138)	0.028 (0.143)
Female × Post × Intensity	0.100 (0.313)	0.395 (0.274)	0.573** (0.286)
Num.Obs.	3150	3150	3121
R2	0.008	0.208	0.250

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: LPM, full sample. Treatment intensity (Pisa faculty, centered). SE clustered at SSD. Controls as in Table 3.10. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.A. REGULATORY HISTORY AND NATIONAL RULES

191

Table A26: Triple difference on winning, both-genders sample (preferred intensity).

	No FE	CUN FE	SSD FE
Female	0.076*	0.028	0.017
	(0.039)	(0.038)	(0.037)
Post reform	−0.007	0.002	0.010
	(0.017)	(0.019)	(0.020)
Treatment intensity	−0.016		
	(0.134)		
Female × Post	−0.004	0.003	0.010
	(0.046)	(0.045)	(0.045)
Female × Intensity	0.153	−0.098	−0.152
	(0.260)	(0.247)	(0.258)
Post × Intensity	−0.343**	−0.107	−0.008
	(0.137)	(0.134)	(0.133)
Female × Post × Intensity	0.357	0.476	0.574*
	(0.303)	(0.291)	(0.293)
Num.Obs.	2195	2195	2195
R2	0.014	0.138	0.170

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: LPM. Sample restricted to competitions with at least one man and one woman. Treatment intensity (Pisa faculty, centered). SE clustered at SSD. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

CHAPTER 3. GENDER-BALANCED COMMITTEES

Table A27: Triple difference, national-faculty intensity (contestable sample).

	No FE	CUN FE	SSD FE
Female	0.043 (0.033)	0.006 (0.032)	0.014 (0.033)
Post reform	−0.011 (0.016)	0.005 (0.015)	−0.008 (0.016)
Treatment intensity	0.004 (0.155)		
Female × Post	0.016 (0.041)	0.011 (0.040)	0.009 (0.041)
Female × Intensity	0.110 (0.264)	−0.153 (0.270)	−0.135 (0.296)
Post × Intensity	−0.196 (0.146)	0.022 (0.146)	0.003 (0.152)
Female × Post × Intensity	0.239 (0.303)	0.376 (0.314)	0.480 (0.330)
Num.Obs.	2965	2965	2965
R2	0.005	0.148	0.180

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: LPM, contestable sample. Treatment intensity based on the national pre-reform female faculty share (centered). SE clustered at SSD. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.A. REGULATORY HISTORY AND NATIONAL RULES

193

Table A28: Triple difference, Pisa-committee intensity (contestable sample).

	No FE	CUN FE	SSD FE
Female	0.034 (0.032)	0.004 (0.031)	0.013 (0.032)
Post reform	−0.011 (0.016)	0.006 (0.016)	−0.006 (0.015)
Treatment intensity	−0.084 (0.149)		
Female × Post	0.022 (0.039)	0.014 (0.037)	0.010 (0.039)
Female × Intensity	0.034 (0.218)	−0.226 (0.228)	−0.203 (0.243)
Post × Intensity	−0.210* (0.127)	−0.027 (0.148)	−0.062 (0.161)
Female × Post × Intensity	0.397* (0.237)	0.474* (0.267)	0.550** (0.268)
Num.Obs.	2965	2965	2965
R2	0.008	0.148	0.181

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: LPM, contestable sample. Treatment intensity based on the pre-reform female share in Pisa committees (centered). SE clustered at SSD. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A29: Triple difference excluding researcher competitions with ≥ 6 candidates, contestable sample (preferred intensity).

	No FE	CUN FE	SSD FE
Female	0.048 (0.036)	0.003 (0.035)	0.008 (0.037)
Post reform	0.014 (0.019)	0.004 (0.019)	-0.019 (0.021)
Treatment intensity	-0.023 (0.146)		
Female $\times$ Post	0.011 (0.046)	0.005 (0.043)	-0.011 (0.046)
Female $\times$ Intensity	0.044 (0.266)	-0.215 (0.247)	-0.246 (0.278)
Post $\times$ Intensity	-0.333* (0.174)	-0.073 (0.141)	-0.053 (0.163)
Female $\times$ Post $\times$ Intensity	0.267 (0.343)	0.523* (0.303)	0.695** (0.329)
Num.Obs.	2071	2071	2071
R2	0.010	0.165	0.205

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: LPM, contestable sample, excluding researcher competitions with six or more candidates. Treatment intensity (Pisa faculty, centered). SE clustered at SSD. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.A. REGULATORY HISTORY AND NATIONAL RULES

195

Table A30: Triple difference excluding researcher competitions with ≥ 6 candidates, both-genders sample (preferred intensity).

	No FE	CUN FE	SSD FE
Female	0.085** (0.041)	0.029 (0.040)	0.014 (0.040)
Post reform	0.007 (0.020)	0.002 (0.022)	-0.009 (0.026)
Treatment intensity	-0.044 (0.162)		
Female $\times$ Post	-0.007 (0.052)	-0.014 (0.051)	-0.007 (0.051)
Female $\times$ Intensity	0.143 (0.286)	-0.136 (0.278)	-0.205 (0.297)
Post $\times$ Intensity	-0.466** (0.186)	-0.162 (0.170)	0.085 (0.195)
Female $\times$ Post $\times$ Intensity	0.370 (0.387)	0.508 (0.357)	0.670* (0.369)
Num.Obs.	1463	1463	1463
R2	0.020	0.150	0.190

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: LPM, both-genders sample, excluding researcher competitions with six or more candidates. Treatment intensity (Pisa faculty, centered). SE clustered at SSD. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A31: Triple difference excluding truncated competitions, national-faculty intensity (contestable sample).

	No FE	CUN FE	SSD FE
Female	0.050 (0.035)	0.005 (0.034)	0.012 (0.037)
Post reform	0.014 (0.020)	0.002 (0.018)	−0.020 (0.020)
Treatment intensity	−0.035 (0.170)		
Female × Post	0.008 (0.046)	0.001 (0.043)	−0.017 (0.046)
Female × Intensity	0.135 (0.290)	−0.153 (0.295)	−0.152 (0.337)
Post × Intensity	−0.289 (0.182)	0.012 (0.172)	0.005 (0.193)
Female × Post × Intensity	0.131 (0.364)	0.332 (0.367)	0.500 (0.409)
Num.Obs.	2071	2071	2071
R2	0.008	0.164	0.204

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: LPM, contestable sample, excluding researcher competitions with six or more candidates. National-faculty intensity (centered). SE clustered at SSD. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.A. REGULATORY HISTORY AND NATIONAL RULES

197

Table A32: Triple difference excluding truncated competitions, Pisa-committee intensity (contestable sample).

	No FE	CUN FE	SSD FE
Female	0.041 (0.034)	0.003 (0.033)	0.012 (0.036)
Post reform	0.013 (0.020)	0.003 (0.018)	−0.018 (0.019)
Treatment intensity	−0.140 (0.161)		
Female × Post	0.017 (0.044)	0.006 (0.041)	−0.016 (0.045)
Female × Intensity	0.067 (0.236)	−0.199 (0.247)	−0.180 (0.268)
Post × Intensity	−0.342** (0.163)	−0.022 (0.184)	−0.074 (0.200)
Female × Post × Intensity	0.462 (0.290)	0.515 (0.322)	0.655* (0.332)
Num.Obs.	2071	2071	2071
R2	0.014	0.165	0.205

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: LPM, contestable sample, excluding researcher competitions with six or more candidates. Pisa-committee intensity (centered). SE clustered at SSD. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A33: Balance test: predetermined characteristics on the triple interaction (contestable sample).

	Log docs	Log cites	Pisa
Female × Post × Intensity	−0.849 (1.033)	−2.312 (1.987)	−0.627* (0.341)
Num.Obs.	2965	2965	2965

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Each column regresses a predetermined characteristic on the triple interaction. Only the triple coefficient is shown. SE clustered at SSD. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

CHAPTER 3. GENDER-BALANCED COMMITTEES

Table A34: Clustering robustness, triple difference (contestable sample, SSD FE).

	SSD cluster	Competition cluster
Female	0.011 (0.034)	0.011 (0.039)
Post reform	−0.005 (0.016)	−0.005 (0.017)
Female × Post	0.012 (0.041)	0.012 (0.043)
Female × Intensity	−0.188 (0.244)	−0.188 (0.292)
Post × Intensity	−0.072 (0.122)	−0.072 (0.130)
Female × Post × Intensity	0.604** (0.275)	0.604* (0.322)
Num.Obs.	2965	2965
R2	0.181	0.181

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$


Notes: Columns differ only in the clustering level. Treatment intensity (Pisa faculty, centered).

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.A. REGULATORY HISTORY AND NATIONAL RULES

199

Figure A2: Leave-one-out: triple-difference coefficient excluding one CUN sector at a time.



Notes: Each point is the triple-difference coefficient (Female $\times$ Post $\times$ Intensity) re-estimated dropping one CUN sector; bars are 95% confidence intervals. The vertical line is the full-sample estimate. Preferred intensity (Pisa faculty, centered), SSD fixed effects, contestable sample.

CHAPTER 3. GENDER-BALANCED COMMITTEES

Titre : Nouveaux résultats empiriques sur la recherche d'emploi

Mots clés : économie du travail, recherche d'emploi, chômage

Résumé : Cette thèse réunit trois essais indépendants d'économie du travail empirique, chacun abordant un problème distinct de la façon dont les marchés du travail allouent et sélectionnent les travailleurs, à partir de terrains très différents.

Le premier chapitre étudie comment une hausse de salaire à l'échelle d'un métier affecte le nombre et la qualité des candidatures ainsi que les embauches. À partir d'une plateforme de travail hôtelier en France, où le salaire par défaut d'un métier a été relevé tandis que les métiers voisins restaient inchangés, il montre que la hausse a fortement accru les candidatures, traduisant une réponse marquée de l'offre de travail, la qualité des candidats n'augmentant que légèrement. Un modèle structurel montre que le report depuis les métiers voisins n'explique qu'une faible part de l'effet. Le deuxième chapitre demande si l'heure à laquelle

se tient un entretien avec un conseiller affecte les trajectoires des demandeurs d'emploi. À partir des registres suisses de l'assurance chômage, il montre que l'heure ne compte que modestement en moyenne, mais de façon très inégale : l'avantage du tôt-matin est bien plus prononcé pour les femmes que pour les hommes. Un choix administratif sans intention distributive façonne ainsi à la fois l'efficacité du service et la répartition de ses bénéfices.

Le troisième chapitre étudie une réforme des comités de recrutement académique dans une université italienne, associant un quota de genre à l'affectation aléatoire des évaluateurs. En exploitant le tirage au sort, il ne trouve aucune preuve que l'équilibre de genre des comités affecte la probabilité qu'une femme soit sélectionnée, excluant des effets importants mais pas plus modestes.

Title : Novel empirical evidence on job search

Keywords : labour economics, job search, unemployment

Abstract : This thesis collects three independent essays in empirical labour economics, each taking up a distinct problem in how labour markets allocate and select workers, and together drawing on very different empirical settings.

The first chapter investigates how an occupation-wide wage increase affects the number and quality of job applications and subsequent hiring. Exploiting data from an online platform for hospitality work in France, where the default wage for one occupation was raised while neighbouring ones were left unchanged, it finds that the higher wage substantially increased applications, pointing to a sizeable labour supply response, while applicant quality rose only modestly. A structural model shows that displacement from neighbouring occupations accounts for only a small part of the effect.

The second chapter asks whether the hour of the day at which a caseworker meeting is scheduled affects the outcomes of unemployed job-seekers. Drawing on Swiss unemployment registers with precisely timed meetings, it finds that timing matters only modestly on average but very unevenly: the early-morning advantage is far more pronounced for women than for men. An administrative choice with no distributive intent thus shapes both the effectiveness of the service and who benefits from it.

The third chapter studies a reform of academic hiring committees at an Italian university that paired a gender quota with the random assignment of evaluators. Exploiting the lottery, it finds no evidence that committee gender balance affects a woman's probability of being selected, ruling out sizeable but not more modest effects.