

An Adaptive Targeted Field Experiment: Job Search Assistance for Refugees in Jordan

Podcast Transcript



Simon Quinn

Welcome to the CSAE Research Podcasts. This is a series of conversations about projects taking place through the Centre for the Study of African Economies at the University of Oxford. I'm Simon Quinn. I'm an associate professor at the Department of Economics and Public Policy at Imperial College Business School and until last year I was privileged to be Deputy Director of the CSAE.

Today we're going to be talking about the project 'An Adaptive, Targeted Field Experiment, Job Search Assistance for Refugees in Jordan'. This was a project that was run in partnership between the CSAE and the International Rescue Committee (IRC).

When we talk about this project, we're talking about what I think is a really important policy question, namely, how is it that different kinds of policy can help refugees and other displaced populations? And we all know, sadly, that we live in a time when this is a really important issue for policy in a lot of different geographical settings. When we talk today about this issue, we're thinking about Syrian refugees and also local jobseekers in Jordan. The labour market in Jordan is characterised by very low employment rates, at least by international standards, and employment rates among refugees are much lower than among Jordanians. So this is a project, as we're going to discuss, in which we think about the impact of three interventions that were designed to improve formal employment outcomes, both for Syrian refugees and for local jobseekers.

I'm really excited to be joined today by two of my co-authors on this project, Stefano Caria, who's a professor in the Department of Economics at the University of Warwick, and Max Casy, a professor in the Department of Economics at the University of Oxford. Max and Stefano, thank you so much for joining me. It's great to be able to discuss this project to you.

Max Casy

Hi Simon. Thanks so much for having me on this podcast.

Stefano Caria

Thank you, Simon. It's really great to be here and to be able to talk about this project.

Simon Quinn

Maybe if I could start with you, Max. I mean, one of the things that I found fascinating working with everyone on this project is that this is an adaptive field experiment. And I learned a lot about this

whole idea of adaptive experiments by working on it. And it's something that I actually think is one of the most exciting parts of this research paper and hopefully something that will be relevant to a bunch of other researchers in this and other field. So can I ask you, first of all, to tell everyone what an adaptive experiment is and how we implemented the adaptive experiment in this context?

Max Casy

So we ran this study in the form of an adaptive experiment and more specifically, an adaptive targeted experiment. Maybe let's back up a little bit and talk about what the purpose of the experiment is or what the goal is, because that really points to why this might be a good idea and why it might be useful in many other experiments.

Over the last couple decades or so in development economics, it has become very standard to run randomised controlled trials. In a randomised controlled trial, you randomly assign different people to the different treatment arms as it's called, different policies, and then you see afterwards how things turn out and you get hopefully a credible estimate of how effective your policies are. That makes perfect sense, I think, for a lot of cases. What that does is it runs the experiment in a way where the goal is to get precise estimates of how effective a policy may be precisely to compare. But that might not be the only thing you care about. Right? So getting precise estimates is not the same thing as informing policy choices, and it's also not the same thing as helping participants. So those are two to rather different objectives. And when we want to pursue this other objective, then we might need to run our experiment in a different way. To illustrate this, a very clear example might be clinical trials in medicine, where you have a drug and a placebo or maybe some standard treatment that you compare to a new treatment. So then you run an experiment, and if you're sticking to a standard protocol, you just have a 50/50 split between the two treatment arms and you observe how well they're doing and you just keep running with that until the end of the experiment. But now what might happen in such an experiment is that halfway through you realise that the new drug is killing everybody, that's not very good. Or it might be that the new drug is a miracle and is curing everybody. And so you would want to use that information for the sake of the patients. You wouldn't just want to keep going with your protocol ignoring what you learned, just use your patient as a means to an end, you want to help your patient. And so what that means is you might want to adapt over time. You might want to shift your treatment or your assignment to treatment or to treatment arms that perform better. So if that new drug is curing everybody then you want to assign the new drug to everybody or as time progresses, if it seems to hurt people, then you might want to quickly shift away. And so that's kind of the core of the activity, and that's something that we implemented now in our experiment, which ran over a whole number of months. And in our experiment, the goal was to help Syrian refugees find jobs in the Jordanian labour market. In order to do that, we observed the different treatments, and Stefano, I guess, can tell us more about this later. The different information and cash interventions to try and help people find a job. Over time the idea was to shift the treatment towards those interventions that were most successful at finding people jobs.

Simon Quinn

Yes, absolutely. Thanks very much, Max. Maybe can I bring you in, Stefano, to tell us in brief about the different treatment arms? And then I want to go back to Max and think about how we implemented this in a way that tries to think about different subgroups. But let me hold off on that until we hear a bit about the treatments themselves.

Stefano Caria

So we had three interventions that we were interested in trying out in this project. One was a small unconditional cash grant, which we calculated to cover the costs of job search for about a couple of months. So we had baseline data on how much people were spending on job search. And so we calculated an amount that people would spend in about a two month period. We know that a job search can actually be surprisingly expensive because people have to use transport and travel around town and have all sorts of expenditures that when you live on a very low income, can actually make up a large share with a surprisingly large share of your budget. And so we try to support them in making these investments by making cash available upfront at the beginning of the search. So cash was the first intervention.

The second intervention was an information intervention where we basically tried to coach people to change people in presenting themselves to employers. Prior research that we've done with you, Simon, and other researchers in Ethiopia, South Africa and Uganda shows that jobseekers ability to convey their employability, their skills and their net worth to employers, it's really a key ingredient to their success in the labour market. And so in the second intervention, jobseekers who may not necessarily be familiar with the new standards of the Jordanian labour market because they were refugees, were basically trained in approaching Jordanian employers and presenting their skills and convincing them of their employability.

And finally we had a psychological or nudge intervention, which was highly motivated by the idea, which is well documented, that Job said it's a very difficult and frustrating process which requires you to sustain the motivation to search hard over a long period of time. So motivation was central to the success of job search, and people struggle to maintain high levels of motivation when they start getting rejection after rejection. So to help people maintain the high level job search effort, we use the insights from recent psychological literature as well as a recent trial in South Africa, and devised a goal setting intervention where jobseekers were encouraged to set themselves some jobs to achieve our goals. And then they would report on whether they attained those goals or not. They would be able to set the other goals for the following week. So these were the three interventions cash, information, and the psychological nudge. And of course, there was also a support group, a control group, that didn't receive any of these three policies.

Simon Quinn

Thank you, Stefano. I remember a fascinating visit that you and I had with some of our partners from the International Rescue Committee, where we actually went from Amman up to the

northwest of Jordan to Irbid and spoke to several refugees. And I think I think we both felt it was fascinating and humbling to hear about the whole constellation of different challenges that these respondents were facing. And I think it's exactly as you say. Partly it's about motivation, partly it's about understanding the labour market, partly it's about having the cash to actually access jobs. And I think that's absolutely right. This is why we went for these three treatments in this context. Max, we opened by talking about targeting. Can I come back to you, having spoken about the treatments, to tell us a little bit more about the idea of targeting and how we implemented it in this setting?

Max Kasy

That sounds good, Simon. Maybe before we get to the targeting, though, let me chat a little bit more about these other activities that I started out with. As Stefano mentioned, we had three different interventions with the control arm, and we talked about this idea, how you want to move over time towards the more successful treatments. This is something that no one is exploiting the information that you have in the machine learning literature. But there is kind of a tension there between exploitation and exploration, which is kind of the other objective, exploration meaning that you keep experimenting in order to actually learn more precisely what is effective. The key challenge when you do this type of adaptive trial is to do three things in balance. You might be greedy and just go all in with something that looked best in the beginning and only do that for the rest of time. But that might be a bad idea because you might just get stuck on something that was just randomly good in the beginning, but actually not such a good idea, or you might be very conservative and slow and just keep experimenting, but then you actually wasting the life chances of your participants. And the key thing is how fast you move from exploring towards exploiting what you've learned as that one thing we put a lot of attention to here.

And then the other thing that you're mentioning is targeting, the idea that everything might not work for everybody in the same way, right? Maybe if you're an older woman with higher education, but no prior labour market experience, who was deployed from Syria, the type of interventions that are useful for you might be quite different than if you're, say, a young man who just dropped out of high school who grew up in Jordan. So it's important to figure out not only what's effective on average, but what effective for whom. And the idea of targeting is that you learn what works for whom, and then actually give people the interventions that make sense for them.

So that's the second innovation or experimental design that we're not only adapting over time towards the more successful interventions, but asking are we using targeting interventions that make sense for different groups of participants. And to do that there's a key challenge, which is how do you combine information, right? So maybe you've never in your data before had an older woman with higher education, but no prior labour market experience or something like that. But then you might draw information from other maybe partially similar participants of the experiment, but then over time you might have more and more information for different groups and then you can just focus on learning from these groups particularly. And so that that again, something that we implemented there, it's a Bayesian Hierarchical Model, but the basic idea is that you combine this information in the optimal way between the group that you're targeting and other groups that

might be similar. So that's in a nutshell, the idea of this experimental design. You're adapting over time to the better interventions, you're targeting the groups for which the interventions make sense. This might sound a little bit complicated or daunting if you're running an experiment yourself. But I want to emphasise, at the end of the day, it's not that complicated. There are a few tools, we provide code and apps, but if you want to run an experiment yourself, the style can just be downloaded and used. In the modern digital era it's actually not too hard to implement an experiment like this. Then you would have in the field offices for unemployed job seekers. They might have some tablets or smartphones where you can just look on a little app to decide which intervention to assign them to. And then in the background we can have a little program running that actually implements our adaptive targeted algorithm. And so that worked out surprisingly smoothly, I would say. And it's actually, I think, quite adaptable to the many different settings.

Simon Quinn

Thanks, Max. I remember very well the discussions in the research team, it was almost like logging in to check the scores of your favourite football team, but especially in the early days, we could actually see the assignment probabilities going up and down on a dashboard of sorts that I think you had kindly prepared. Can you tell us just a little bit more? Because I think the intuition is clear here, but the algorithm is handling all of the assignment probabilities, and then the algorithm is going to give assignment probabilities. And then what? And then the researcher then uses those to randomise or indeed the code then just generates a randomisation for somebody else who enters the program the next day, as it were. Is that a fair intuitive summary?

Max Kasy

The dashboard was the back end for us to follow what was going and that was interesting. On the level of the people in the field who were actually implementing this, they don't have to deal with any of that. They just go to the website or app, where they can enter that they have a jobseeker with these characteristics and that gives them like the information dimension or gives them some cash to help them search for a job. Everything else is kind of run in the background, but the program that collects the data over time, how successful different interventions were, that's collected in a spreadsheet somewhere and then the program without the spreadsheet decides how to assign participants of the program to different treatment arms.

Simon Quinn

One of the interesting things about doing this in a labour market context is that you don't necessarily have to wait six, 12, 18 months to know whether a policy is or isn't working. In this context, as you alluded to, we were able to have relatively quick return information on whether someone had found a job in the short term. And of course, that may mask longer term effects. Some treatments in some contexts may take a long time to materialise. But I guess the point is, in

this case, we thought that that was going to be a pretty good proxy for people's longer term success on different treatments.

Max Kasy

That's exactly right. So participants got a call, I think six weeks or two months after participating in the program, and then they were just asked on the phone if they'd found a job. And that's the key outcome that we're trying to maximise, the probability of people finding a job a couple of months after participating.

And that might be quite easy to implement in some settings and not easy to implement in other settings, right? So this type of method is used a lot in online settings, like when the big tech companies are maximising advert clicks, that's something where you get very quick feedback. Did somebody click on this link that we put on the website or not? In that kind of situation it's very easy to adapt quickly and efficiently. But then there might be other settings where it takes a long time for problems to realise. A great example, I think, typical interventions that would affect Alzheimer risk, it might take 30, 40 years for any results to be seen. And so in that kind of setting it would not be very useful to think about adaptive methods because you won't have outcomes that quickly. But a lot of the types of policies you might be interested in in the labour market or development economics contexts, it would not be unexpected to have something quite effective within the timeframe of an experiment. And as long as you have observed outcomes for some people before the end of the experiment, then you'll have information that you can use in order to take better care of your participants. And again, I want to emphasise the ethical dimension of that, not just seeing them as a means to an end, but seeing the participants as an end in themselves. Probably not just using them to learn for our academic papers or our policy recommendations, but we're really trying to do right by them, by giving them the interventions that are most helpful to them.

Stefano Caria

One thing that I would add for the people who are listening is that one area where I expect this to be particularly applicable is with experiments that are running with large institutions that have a constant flow of clients, for example, schools, hospitals or courts or things like that, where you basically have a fixed institution that serves a population that gets the service and then leaves, and you need a lot of them. We do a lot of work with these kind of institutions. And in this case you get both fast feedback and a constant stream of new people that would benefit from the adaptivity. So I just wanted to flag that there are many settings where we work as applied economists where this would actually be a good match.

Simon Quinn

I think that's a great point. I think also fascinating that you raise the point of the partner institution. I think this dovetails nicely with Max's point about ethics. I would find it really difficult to go into a crisis situation where you've got, for example, a large non-government organisation (NGO)

trying to help a target population and say, “Hi, can I do a long term fixed proportion randomised controlled trial (RCT) and please don't do anything with the control group for two years because I need to write my paper” I don't think any researcher would feel comfortable with that. So I think this is where the ethical dimension that Max raises touches on the institutional dimension that you can say to a policymaker, look, we're going to experiment, but we're going to do our very best to move respondents towards treatments that seem to be helping them and not just helping on average, but, as Max explained, helping for the particular kind of respondent that they are.

Max Kasy

I think it's a great point. In particular, oftentimes, I think when you the RCTs in a more classic way, then you have this tension between an implementation partner organisation and the researchers, The implementation partners often believe that they know what works best to help their clients, their goal is not to do research. And then you have the researchers who want to write academic papers with precisely estimated treatment effect. And those are two different goals. And so I think adaptive experiments, what we are contributing here, what is allows you to do is to align more closely the objective of the experimenters with the objectives of implementation partners.

In our case, the International Rescue Committee said their goal is to help refugees find job quickly, and we could say that's our goal too and we're designing our experiment to help refugees find jobs quickly There is no conflict here, and so you should be happy to let us do the experiment because we are actually just helping you to fulfil your mission as opposed to kind of being in a tension with what you're trying to achieve. And so that might actually also open all kinds of new sites and venues where you could run experiments where a classic RCT might be unethical but also politically unfeasible.

Simon Quinn

I completely agree. In a minute, I want to ask Stefano to actually tell us about the results. Before I do that, let me pose one more question to you, Max. And I might I might make friends and enemies in the world of Bayesian statistics here, but I want to throw in a slightly nerdy question. Suppose you're dealing with a policymaker who says, Alright, I'll do this adaptive RCT, but I have a very strong belief that treatment one is going to work and treatment two is not. Some people would say you should use that prior belief to seed the model. So are we going to start by putting a lot more people into the treatment arm, reflective not of data that we've collected in the experiment, but reflecting the strong prior that the policymaker has. Other people would say, no, it's interesting and important that the policymaker feels they know what works. But we should start with a much less informative prior, to use the terminology, and let the experiment in the data tell us what's working.

I know there are different views out there about this. In our context, we collected information about what policymakers thought, but we didn't feed that information directly into the algorithm.

Max I'm just keen to draw you out just for a few more minutes on your thoughts about the respective pros and cons of those approaches.

Max Kasy

So I think that that makes perfect sense for policymakers, partners and clients for that matter, often have a lot of information about what might or might not work. Right? So we learned that in our experiment from running some focus groups with refugees who had very strong opinions about the interventions which largely lined up with what we found later in the experiment. And so that kind of information can be very useful for deciding to experiment. Let's just say, if we have a strong belief beforehand that some intervention will be more effective then it makes a lot of sense that that intervention to get at a larger sphere for the observation or for our participants to get assigned to that intervention. I think it's very important to think separate here. One is how you run the experiment. You might put a lot of people in the information intervention and less people in the nudge intervention or something like that, but that doesn't prevent you from analysing the experiment, after the fact, in a way that your analysis doesn't draw on that prior at all. Right? So if you just want to have the get scientific evaluation, that only depends on the data and does not draw on the policymaker or a partner or a client beliefs, that's perfectly fine. You can do that even if your experiment where you used those policymaker or partner beliefs in order to decide which treatment is assigned more often.

Simon Quinn

Understood. Thanks, Max. I feel like we've been teasing our listeners for a while now with the discussion of the actual results. Let me get back to you, Stefano. Why don't you tell us about what we think we learned in this context, about the constraints that refugees face in the labour markets and how we think that might be relevant for other kinds of policy interventions in similar settings in the future.

Stefano Caria

Thank you. I'll talk about the treatment effects, but perhaps let me say a couple of words about the population that we are working with just to set this in context. So our key population was a population of Syrian refugees found in three cities around Jordan. These are the population that perhaps in Syria was not necessarily particularly poor, but wasn't doing very well in Jordan. And so employment rates were extremely low and people were very close to the poverty line, many of them below the poverty line. And there was quite a bit of interesting work and quite a bit of work experience as well, and but very little at least formal employment. We know that some people were working informally, although they were not necessarily very comfortable saying that as this would break the law. And so given that they are quite vulnerable as refugees, this was something that they weren't really forthcoming in discussing with us. Now, we then offered these three interventions. We offered these interventions both to the Syrian refugees, that were the central population of our study, but also to a second population of local Jordanians. This is actually quite important for two reasons. One, because we want to see whether the same interventions within the constraints that prevent these two population from accessing the labour markets are similar or whether they need different policies.

And two, going back to ethics, it's a very important commitment of the IRC to always match the support that they give to the population of a given country that is hosting to an equivalent amount of support given to other national individuals living in the same areas that are also struggling to find employment. So it's a kind of requisite for IRC to do this kind of equitable policy allocation to avoid resentment in the perception that the refugees are favoured. So I would say we also can bring a second population of national Jordanians. Now, what did we find? We interviewed individuals six weeks after treatment and two and four months after treatment. So we have three points in time in which we can observe outcomes. Six weeks after treatment, it was too soon, we found that there were minimal effects on employment for all groups. Two and four months after treatment, we found that the cash intervention in particular, increase the job search and then in turn employment and earnings. Now, the type of treatment effects that we are talking about, this is close to four percentage point increase in employment, which we should put in context in at least two ways. Number one, four percentage points, it's not huge. This is clearly not transformational in a sense it's a light touch intervention, so perhaps we shouldn't expect them to transformation effects. We are basically getting the employment rates from about 9% to about 13%. So this is still a population that is fairly locked out of the labour market. However, in relative terms, this is a fairly large increase in employment. So it's a 50% increase. And also relative to the kind of treatment effects that we found in the actual labour market policy literature, the research literature that they decide to devise policies in other contexts in the world. It's actually in line with what people found elsewhere. In fact, it's a towards the upper part of the distribution of treatment effects. Same story for the treatment effects of income. We found that income goes up by almost 60%, the absolute amount, it's limited in a sense because we're starting from a very low base, but in relative terms, this is a large effect.

One more thing that I want to say about the cash interventions, that this is one of the few unconditional cash interventions that are designed to boost search. There are other papers in this literature that show that if you gave people conditional cash, for example, such as vouchers that enable people to take transport to the centre of town. Those interventions boosted job search. But it wasn't clear whether that was the case because you were almost forcing people to search for work in order to benefit from the intervention or whether people would have increased their job search, given some additional resources, whether this was something that was very in their high priorities. They were not doing extra investment in job search because of their lack of cash. So it's quite interesting that we find that giving these unconditional cash grants, amongst Syrians, does raise job search quite a bit. And then this leads to employment.

Among Jordanians, on the other hand, while this is still a fairly poor population, but they're richer than the Syrians, they find jobs at faster rates and are generally better off. We actually find the cash has very limited effects. The other two interventions, the information and nudge interventions also increased job search, less so than the cash intervention. And they do lead to some employment effects at the two month marker. But these effects are smaller and less persistent. So by the time of the four month interview, we see positive coefficients but smaller and in general statistically significant.

So these are the two interventions, although they seem to have at least activated some job search and generated some gains, they proved less effective.

Simon Quinn

Thanks Stefano, I mean, it's a really interesting summary. This issue of how you help displaced people in urban labour markets is sadly not an issue that's going away. So the takeaway for a policy setting for future crises is that, when we think about cash support for refugees or displaced populations, that this is not only valuable intrinsically in terms of increasing welfare, but we should actually think of this as unlocking search capacity in potentially new labour markets. Is that a fair takeaway? I know as researchers we always struggle to think about how to generalise our results to other contexts, but it feels to me like that's a sort of insight that we gained from this population that may be relevant for future crises.

Stefano Caria

Absolutely. So I would say that basically one important takeaway, of course if you generalise to different crises, other things will change. But it's probably true overall that refugees typically lose a lot of assets during the process of displacement. So this is a population where a little bit of cash can have very high returns. The point stands that also in our population, employment rates and incomes remain quite low. So it's not it's not going to be by any means a sufficient intervention, but it's probably a small but highly targeted thing that one can do to help refugees operate in the new labour markets. That is one more thing that I want to say, which is of course there are many things that would differ when you compare this to refugee crisis. But I think one factor that is of essential importance is language. And so some refugee crises, for example the one experienced by Jordan or Lebanon, you find that refugees coming to a country speaking effectively, the language that is spoken in their country, maybe a different dialect, maybe they have a bit of an accent, but by and large they're already able to communicate with pretty much anybody in the host country. The refugees are going to Turkey, for example, in a very different position. In fact, in Turkey we find that they tend to mostly work in self-employment and keep trading with Syrians living elsewhere in the world rather than being getting integrated in the local economy like the Jordanian and Lebanese policymakers were trying to do. I'm not entirely sure which of the two kinds of shocks is more common. For example, Venezuelans going to Colombia, and they spoke the same language, a different accent but the same language. Lots of other refugees and some of those in the in Eastern Africa, for example, you find the populations that that are really, and many refugees base in Europe, language is the first barrier. So I would say that in cases where language is also an issue, I'm not entirely sure this is also generalised simply because refugees may not be immediately employable before they pick up the essential language skills.

Simon Quinn

That's a really excellent point. Let me loop back to a point that Max made earlier, which is that we did a bunch of qualitative work, of course, talking to this population and this narrative about the importance of cash for unlocking search was certainly one that seemed to resonate with a lot of the

refugee respondents themselves, which I think is interesting, and maybe also something that would be useful in terms of learning in other contexts. Exactly. About the practicalities that you talk about.

Well, I've really enjoyed this conversation. It's brought back a lot of fond memories of a project that I really enjoyed doing and where I certainly learned a lot.

Thanks very much, Stefano and Max, and thanks to all of our listeners for tuning in. I'm certainly very keen to see how these two different literatures go to see where people think about new interventions to help these kind of vulnerable displaced populations and also to see the emerging field of adaptive and targeted field experiments. Thanks, everyone.

Stefano Caria

Very exciting. Likewise. Thank you so much.

Max Kasy

Thank you so much. Thanks everyone.