Preface

The LSE SU Economics Society's Research Division is incredibly proud to present the third volume of Rationale, LSE's only student economics journal.

Rationale was first established as a working paper series in 2019, by Christopher Dann, the Society's Head of Research in 2018-2019. In 2020, Rationale became the first student economics journal at LSE under the leadership of Kerry Neitzel and Xuyi Yang.

This is the second year we publish Rationale as a student managed and edited research journal, now affiliated with the Houghton St Press. This is also the first year where all of the Division's work has gone online, given challenges due to COVID-19. Despite this, the division has grown more collaborative from the start, with work-in-progress seminars, internal presentations and the final public presentation, where student researchers present their year-long work.

We are also proud to have launched the first student economics research incubator programme, an extension of the Lightning Research programme developed by former Economics Society members. In three weeks, we provided interested students in various disciplines with a quick training and preview of quantitative research in economics.

We would like to express our sincerest gratitude to this year's working group leaders Jason Jia, Stefanus Phan, Mina Burns, Merton Ngan and Eddy Zou, for their determination and efforts to lead working groups against numerous challenges and supporting each other with valuable suggestions. The research projects were only enabled by the hard work of the Research Associates, who are incredibly resilient to conduct research remotely, thank you for your commitment and work. Thank you also to the Editorial Board for providing helpful feedback and comments to the working groups. Special thanks go to Marco Canal who has been a co-editor at Rationale for two years.

The working groups benefited greatly from the support from professors and faculty in the Economics Department, including Steve Pischke, Rachael Meager, Guy Michaels, Nico Rosetti and Canh Thien Dang, who we thank on behalf of the entire Research Division. We are especially grateful to Judith Shapiro for her invaluable support, advice, thoughtful discussions and suggestions throughout the year.

We also thank the LSE SU Economics Society's executive committee for their support with Research Division activities. Among others, we thank Celine Mano, Sally Yang and Sarah Wang for their valuable inputs and assistance. We also thank previous executive committee members at the Society, especially Samantha Ong, Aaron Luke, Chris Shaw, Tom Glinnan and others, for their valuable assistance with the research incubator programme.

We genuinely hope that the Research Division will further thrive as a wonderful platform for students to learn about research practice and enjoy pursuing their own ideas together with others. Last but not least, we thank the incoming Head of Research, Marie Ogino, for helping make this possible and wish her the best of luck for the year ahead.

We hope you enjoy this year's edition of Rationale!

Eddy Zou, Head of Research and Editor in Chief 2020/21
Mikael Johansson, Deputy Head of Research and Editor 2020/21

Inflation under FIRE: Deviations and Explanations*

Jason Jia, Birce Akay, Patrick Leitloff and Annie Wang†

Apr 2021

Abstract

We examine established empirical evidence on deviations of inflation expectations from the Full-Information Rational Expectations (FIRE) hypothesis. However, these empirical facts offer seemingly contradictory conclusions. We seek to reconcile these findings using a new piece of empirical evidence and a parsimonious model, adapted from the methodology of Angeletos, Huo and Sastry (2020). We confirm that all empirical facts can be explained by a combination of dispersed, noisy information and over-extrapolation.

*We are grateful to Nika Vahčić for her valuable support and assistance in early stages of the project. We thank Eddy Zou and Mikael Johansson for their leadership and continued support. We also thank members of LSE SU Economics Society's Research Division for helpful comments.

†The authors are from The London School of Economics and Political Science.

1 Introduction

The Full-Information Rational Expectations (FIRE) hypothesis, an essential building block of most modern macroeconomic models, assumes that all data about the state of the economy is common knowledge (FI), and agents form the same optimal beliefs about the future (RE).

Often modelled under FIRE, inflation expectations is a macroeconomic variable with key implications for pricing decisions (Coibion, Gorodnichenko and Ropele, 2020), the state of the economy as characterized by the New Keynesian Phillips Curve (Clarida, Gali and Gertler, 2000), optimal monetary policy (Clarida, Gali and Gertler, 1999) as well as asset prices (Stock and Watson, 2003). There are thus huge incentives for firms, economists, central banks and financial agents to form accurate inflation forecasts.

However, a rapidly growing body of empirical (Section 1.1 and 1.2) and theoretical (Section 1.3) work has built a case for systematic deviations of inflation forecasts from FIRE, in the form of over/under-reaction and disagreement. While these facts use the same US data sources (Section 2), they are generated by different sets of authors with somewhat different procedures. Our contribution here is thus to use the common methodology of Bordalo, Gennaioli, Ma and Shleifer (2020) to verify all of them (Section 3). Recently, Angeletos, Huo and Sastry (2020) proposed a new candidate empirical fact using Impulse Response Functions (Section 1.4); we contribute by evaluating the findings using an alternative methodology, in addition to the authors' approach (Section 4). We then build a model based on Angeletos, Huo and Sastry (2020), as well as Reis (2020), to characterize deviations from FIRE using key parameters, and in particular, perceived parameters being different from true parameters. By matching model predictions to the empirical facts and inferring the matching set of parameters, we can eliminate inconsistent theories and identify plausible ones. We subsequently confirm that all empirical facts about over/under-reaction can be explained by a combination of information frictions or dispersed, noisy information (under-reaction) and over-extrapolation (over-reaction); this provides a potential unified explanation (Section 5). Finally, we conclude with key findings, strengths and caveats to note, and ideas for future work (Section 6).

1.1 Established Empirical Facts about Over/Under-Reaction

The growing literature has established 3 empirical facts regarding over/under-reaction (OU) at the aggregate and individual level. However, as seen below, the 3 facts offer inconsistent patterns of under- and over-reaction, and a unifying explanation remains elusive.

OU1. Aggregate forecast errors for inflation are positively correlated to lagged aggregate forecast revisions, as observed in Coibion and Gorodnichenko (2015). This suggests that aggregate forecasts **under-react** to aggregate revisions, which are based on news about actual inflation.

OU2. Individual forecast errors for inflation are negatively correlated to lagged individual forecast revisions, as observed in Bordalo, Gennaioli, Ma, and Shleifer (2020). This suggests that individual forecasts **over-react** to individual revisions, which are based on news about actual inflation.

OU3. Aggregate forecast errors are positively correlated with actual levels of inflation, as observed in Kohlhas and Walther (2018). This suggests that aggregate forecasts **over-react** to news about actual inflation.

1.2 Established Empirical Facts about Disagreement

Under FIRE, no disagreement is predicted, as people share the same information and form expectations conditional on that information. However, another strand of the growing literature has shown this to be counterfactual and established important empirical facts regarding disagreement (D), one of which[1] is summarized as below:

D1. Disagreement in inflation expectations is substantial, varies over time, and increases temporarily with (both positive and negative) shocks. This is observed by Mankiw and Reis (2002) as well as Mankiw, Reis and Wolfers (2003).

1.3 Theories Predicting Deviations from FIRE

At the same time, the theoretical counterparts saw an explosion of alternative theories that seek to explain particular deviations from the FIRE hypothesis, with a selected, non-representative summary presented below. However, there has been little attempt at identifying models consistent with *all* empirical facts.

1.3.1 Deviations from FI (Information Frictions)

Dispersed, Noisy Information. Instead of observing inflation directly, forecasters observe a signal contaminated with white noise (Woodford, 2003; Angeletos and Lian, 2016). Forecasts for inflation today are formed by taking an average of today's signal and yesterday's forecast of inflation today.

Sticky Information. Each period, a fraction θ of forecasters update expectations using RE, while $(1 - \theta)$ do not update expectations (Mankiw and Reis, 2002; Kiley, 2007).

Rational Inattention. Acquiring information is costly, and thus forecasters may rationally make decisions based on incomplete information, if it is less costly than acquiring complete information (Sims, 2003; Mackowiak, Matejka and Wiederholt, 2018).

At the aggregate level, these models all imply that forecast revisions predict forecast errors; that there is under-reaction due to 'stale' expectations. At the individual level, however, RE implies that forecast revisions cannot predict forecast errors.

[1] Two other empirical facts about disagreement include: (1) Policy communication lowers the disagreement that results from a shock, and (2) Regime changes that raise transparency can permanently lower disagreement. These findings are important, but not the focus of our research. It is noteworthy that the final model described is also able to explain these facts, as per Reis (2020).

1.3.2 Deviations from RE (Imperfect Rationality)

Models explaining imperfect rationality can be classified by whether they predict over- or under-extrapolation from available data.

Overestimation of Persistence. The true process follows some AR process with persistence ρ, but forecasters overestimate the persistence, such that $\hat{\rho} > \rho$ (Greenwood and Schleifer, 2014; Gennaioli, Ma, and Shleifer, 2016). This predicts over-extrapolation and thus over-reaction at the aggregate and individual level.

Over-Confidence. Forecasters are over-confident when they believe the precision τ of their private information is higher than it actually is, such that $\hat{\tau} > \tau$ (Broer and Kohlhas, 2018). This also predicts over-extrapolation and thus over-reaction at the aggregate and individual level.

Level-K Thinking. Agents perfectly observe inflation, but assume everyone else is a 'level' lower than them, and that a 'level-0' player always plays a default action (e.g. set forecast = 0); agents then play the best response accordingly (Gabaix, 2020; García-Schmidt and Woodford, 2019). This predicts sluggish convergence, under-extrapolation and under-reaction at the aggregate and individual level.

1.4 New Candidate Empirical Fact

In response, Angeletos, Huo and Sastry (2020) proposed a new candidate empirical fact, which when used with the model developed in Section 5, identifies a possible unified explanation for all empirical facts:

OU4. There is **initial under-reaction** and **delayed over-reaction** in the response of aggregate forecast errors to 'cost-push', or inflationary shocks.

2 Data

We use US data from 1968 Q4 to 2017 Q4 for all variables of interest: inflation forecasts, inflation figures and other macroeconomic variables.

2.1 Forecasts

Our main data source on inflation forecasts comes from the Survey of Professional Forecasters (SPF), issued by the Federal Reserve Bank of Philadelphia. The SPF is a panel survey of about 40 experts from industry, government and academia. Each quarter, forecasters report level forecasts of GDP price index (PGDP) in the current and next four quarters, from which inflation forecasts can be computed.

We will be using the following data sets, which contain quarterly time-series data concerning PGDP: (1) Median aggregate PGDP forecasts in levels, (2) Median aggregate PGDP forecasts in growth, and (3) Individual PGDP forecasts in levels.

For aggregate forecasts, we use the median inflation rate to mitigate concerns about outliers and give a better representation of the 'average' figure. For individual-level data, we remove outliers (which can be extreme) by computing the median and inter-quartile range (IQR), and trimming observations in individual forecasts which are ≥ 5 times the IQR from the median, as per Bordalo, Gennaioli, Ma, and Shleifer (2020).

2.2 Inflation

We use 2 different notions of inflation in our study for 2 different purposes: (1) Inflation (GDP Deflator) and (2) Vintage data in levels.

The conventional inflation measure is found in the Federal Reserve Economic Data (FRED), published by the Federal Reserve Bank of St. Louis. We use this measure for more general purposes, specifically building a SVAR.

As inflation figures are updated across time, however, estimates of inflation between $[t, T]$ can often differ at time S and $S+i$, where $S \geq T > t$ and $i > 0$. This turns important when we look to estimating how inflation expectations change upon release of the most recent inflation figure. For this purpose, we use vintage data from the Philadelphia Fed, which provides the full time path of inflation figures.

2.3 Other Macroeconomic Variables

We will also be using data consisting of quarterly observations of the following 9 macroeconomic variables, published under the FRED by the St. Louis Fed to setup our SVAR model: (1) Unemployment Rate; (2) Real GDP per capita; (3) Real Investment per capita; (4) Real Consumption per capita; (5) Hours worked per person; (6) Level of utilization-adjusted Total Factor Productivity (TFP); (7) Labour Productivity in the Nonfarm Business Sector; (8) Labour Share; and (9) Nominal Interest Rate (Federal Funds Rate).

3 Established Empirical Facts

In this section, we examine key facts documented in existing literature, each presented by different sets of authors. Our contribution here is to increase the robustness of these findings, by verifying all of them using the methodology of Bordalo, Gennaioli, Ma and Shleifer (2020).

As this paper is heavily involved with inflation and inflation expectations, we first develop a more rigorous and unambiguous set of notation than would be normally observed.

Let $\pi_{t-1,t+h,k,l}$ denote the h+1 period inflation from time t-1 to t+h, given as a rate with frequency of $1/k$ year (k=1 for annual rate, k=4 for quarterly rate) with compounding frequency of $1/l$ year (l=1 for annual compounding, l=4 for quarterly compounding). Next, let $\mathbb{E}_t(\pi_{t-1,t+h,k,l})_i$

be the individual expectation (for SPF forecaster i) of h+1 period inflation from t-1 to t+h, in period t. Then, let $\bar{\mathbb{E}}_t(\pi_{t-1,t+h,k,l})$ be the aggregate median expectation (across all SPF forecasters) of h+1 period inflation from t-1 to t+h, in period t. Similarly, let $\bar{\mathbb{E}}_{t-1}(\pi_{t-1,t+h,k,l})$ be the median expectation of h+1 period inflation from t-1 to t+h, in period t-1.

We focus on 4-period inflation, with horizon h=3, annual rate $k = 1$ and annual compounding $l = 1$. To link our desired object to actual SPF data, we compute the median expectations as follows given quarterly rates with annual compounding,

$$\begin{aligned}\bar{\mathbb{E}}_t[\pi_{t-1,t+3,1,1}] &= [\prod_{s=0}^{3} \bar{\mathbb{E}}_t(1 + \pi_{t+s-1,t+s,4,1})^{1/4}] - 1 \\ &= \bar{\mathbb{E}}_t[(1 + \pi_{t-1,t,4,1})(1 + \pi_{t,t+1,4,1})(1 + \pi_{t+1,t+2,4,1})(1 + \pi_{t+2,t+3,4,1})]^{1/4} - 1\end{aligned} \tag{1}$$

and as follows given quarterly rates with quarterly compounding.

$$\begin{aligned}\bar{\mathbb{E}}_t[\pi_{t-1,t+3,1,1}] &= [\prod_{s=0}^{3} \bar{\mathbb{E}}_t(1 + \pi_{t+s-1,t+s,4,4})] - 1 \\ &= \bar{\mathbb{E}}_t[(1 + \pi_{t-1,t,4,4})(1 + \pi_{t,t+1,4,4})(1 + \pi_{t+1,t+2,4,4})(1 + \pi_{t+2,t+3,4,4})] - 1\end{aligned} \tag{2}$$

The aggregate forecast error at time t is defined as follows,

$$Error_t = \pi_{t-1,t+3,1,1} - \bar{\mathbb{E}}_t[\pi_{t-1,t+3,1,1}] \tag{3}$$

and the aggregate forecast revision at time t is defined as follows.

$$Revision_t = \bar{\mathbb{E}}_t[\pi_{t-1,t+3,1,1}] - \bar{\mathbb{E}}_{t-1}[\pi_{t-1,t+3,1,1}] \tag{4}$$

For individual forecast error ($Error_{i,t}$) and revision ($Revision_{i,t}$), we use $\mathbb{E}_k[\pi_{t-1,t+3,1,1}]_i$ instead of $\bar{\mathbb{E}}_k[\pi_{t-1,t+3,1,1}]$ for $k = \{t-1, t\}$.

3.1 OU1: Under-reaction in aggregate forecasts

Coibion and Gorodnichenko (2015) test for deviations from FIRE by running a regression of aggregate forecast errors against aggregate forecast revisions.

$$Error_t = \alpha + \beta_{CG} \cdot Revision_t + u_t \tag{5}$$

Our estimates of β_{CG} are given in Table 1 below, with reported Robust Standard Errors (RSEs). We run regressions for the entire sample from 1968-2017, and also separate regressions for the periods 1968-1983 and 1984-2017. This is due to high inflation volatility prior to 1984[2], and significantly lower inflation volatility from 1984 onwards, largely credited to the increase in central

[2] In particular, the oil price shocks in the 1970s and the Volcker disinflation period from the 1970s till 1983.

bank independence, and the adoption of flexible inflation targeting with reference to the Taylor Rule. In fact, the mid-1980s to 2007 period is often noted as the 'Great Moderation' due to stable macroeconomic conditions. As such, the restricted post-1984 dataset should reflect a more 'stationary' macro environment, and be subject less to the Lucas critique.

Table 1: Regression of Aggregate Errors on Aggregate Revisions

	Dependent variable:		
	$Error_t$		
	1968-2017	1968-1983	1984-2017
	(1)	(2)	(3)
$Revision_t$	1.461***	1.615***	0.384
	(0.256)	(0.267)	(0.290)
Constant	−0.001	0.002	−0.003***
	(0.001)	(0.002)	(0.001)
Observations	187	55	132
R^2	0.301	0.367	0.024
Note:	*p<0.1; **p<0.05; ***p<0.01 (RSEs)		

We find in all specifications a point estimate of $\beta_{CG} > 0$, consistent with original findings. This suggests that aggregate forecasts under-react to aggregate revisions. The coefficient is considerably lower on the restricted post-1984 sample, plausibly reflecting smaller deviations from FIRE in an environment with greater macroeconomic stability.

An intuition behind this link between $\beta > 0$ and under-reaction can be drawn with the support of Figure 1, a combination of three tables.

In the first table, note that a 4-period inflation forecast from t-1 to t+3 (i.e. $\bar{\mathbb{E}}_t[\pi_{t-1,t+3,1,1}]$) requires the estimates of prices for 5 time periods, as highlighted in blue: $\{t-1, t, ..., t+3\}$. Also note that this forecast can be made at time t, at time t-1, or any time before all 5 periods' prices are known.

In the second table, the link between the four consecutive 1-period (quarterly) inflation forecasts ($\bar{\mathbb{E}}_t[\pi_{t+s-1,t+s,4,1}]$ for $s = \{0, 1, 2, 3\}$, in blue) and the single 4-period (yearly) inflation forecast ($\bar{\mathbb{E}}_t[\pi_{t-1,t+3,1,1}]$, in yellow) is shown. Note that 5 time periods give 4 quarterly inflation figures.

In the third table, we construct a period of rising actual inflation, with forecasts at time t increasing at a slower pace given under-reaction; forecast errors thus 'accumulate' and increase over time. Since forecasts eventually align with the truth, we also expect revisions to increase over time. Thus, we expect errors to be positively correlated with revisions. As such, a point estimate of $\beta_{CG} > 0$ suggests under-reaction at the aggregate level.

Time of forecast	Time	t-2	t-1	t	t+1	t+2	t+3	
	E.g.	1969 Q3	1969 Q4	1970 Q1	1970 Q2	1970 Q3	1970 Q4	
t-1 (1969 Q4)			X	O	O	O	O	
t (1970 Q1)			O	X	O	O	O	

Time of forecast	Time	t-2	t-1	t	t+1	t+2	t+3	
	E.g.	1969 Q3	1969 Q4	1970 Q1	1970 Q2	1970 Q3	1970 Q4	4-period
t-1 (1969 Q4)				2%	2%	2%	2%	2%
t (1970 Q1)				2.5%	2.5%	2.5%	2.5%	2.5%

Summary	t	1969 Q3	1969 Q4	1970 Q1	1970 Q2	1970 Q3	1970 Q4	
Actual		2%	2.5%	3%	3.1%	4%	4.5%	Diagnosis: **Under-Reaction**
Forecast at t-1		2%	2%	2%	2.5%	2.6%	2.9%	
Forecast at t		2%	2%	2.5%	2.6%	2.9%	3.5%	
Revision		0pp	0pp	0.5pp	0.1pp	0.3pp	0.6pp	Correlation: **+**
Error at t		0pp	0.5pp	0.5pp	0.5pp	1.1pp	1.0pp	

Figure 1: Figure illustrating the link between under-reaction and a positive coefficient in a regression of errors against forecasts, in a constructed period of rising actual inflation. Reversing the story reveals over-reaction can be captured by a negative coefficient. t refers to different periods as real time progresses. X refers to the current period (i.e. time of forecast), and O refers to periods for which the 4-period measure of inflation takes prices from (if not already captured by X). Inflation numbers in the second table in blue are given as $\bar{\mathbb{E}}_t[\pi_{t+s-1,t+s,4,1}]$ for $s = \{0, 1, 2, 3\}$, while inflation numbers in the third table (including those in the second table in yellow) are given as $\bar{\mathbb{E}}_t[\pi_{t-1,t+3,1,1}]$. The stated inflation numbers are not actual figures.

3.2 OU2: Over-reaction in individual forecasts

Bordalo, Gennaioli, Ma and Shleifer (2020) test for deviations from FIRE by running a regression of individual forecast errors against individual forecast revisions.

$$Error_{i,t} = \alpha + \beta_{BGMS} \cdot Revision_{i,t} + u_{i,t} \tag{6}$$

Our estimates of β_{BGMS} are given in Table 2 below, with reported RSEs.

Table 2: Regression of Individual Errors on Individual Revisions

	Dependent variable:		
	$Error_{i,t}$		
	1968-2017	1968-1983	1984-2017
	(1)	(2)	(3)
$Revision_{i,t}$	0.155^{***}	0.218^{***}	-0.284^{***}
	(0.041)	(0.056)	(0.031)
Constant	0.000	0.006^{***}	-0.003^{***}
	(0.000)	(0.001)	(0.000)
Observations	5,153	1,619	3,534
R^2	0.008	0.015	0.038
Note:	*p<0.1; **p<0.05; ***p<0.01 (RSEs)		

We find that $\beta_{BGMS} > 0$ over the full sample as well as the pre-1984 period, but $\beta_{BGMS} < 0$ in the post-1984 period; this is consistent with original findings. Given the 'more stationary' environment post-1984, we interpret the findings to be supportive of over-reaction at the individual level.

We note that this result is puzzling in light of **OU1**, because forecasters over-react at the individual level, but under-react at the aggregate level.

3.3 OU3: Over-reaction in aggregate forecasts

Kohlhas and Walther (2018) test for deviations from FIRE by running a regression of aggregate forecast errors against current quarterly realizations of inflation.

$$Error_t = \alpha + \beta_{KW} \cdot \pi_{t-1,t,4,4} + u_t \tag{7}$$

Our estimates of β_{KW} are given in Table 3 below, with reported RSEs.

Table 3: Regression of Aggregate Errors on Current Realizations of Inflation

	Dependent variable:		
	$Error_t$		
	1968-2017	1968-1983	1984-2017
	(1)	(2)	(3)
$\pi_{t-1,t,4,4}$	0.019	-0.355^{***}	-0.218^{***}
	(0.043)	(0.106)	(0.047)
Constant	-0.001	0.026^{***}	0.003^{*}
	(0.001)	(0.007)	(0.001)
Observations	187	55	132
R^2	0.001	0.129	0.122
Note:	*p<0.1; **p<0.05; ***p<0.01 (RSEs)		

We find that $\beta_{KW} > 0$ over the full sample as well as the pre-1984 period, but $\beta_{KW} < 0$ in both the pre-1984 and post-1984 period.[3] Given the 'more stationary' environment post-1984, we also interpret the findings to be supportive of over-reaction at the aggregate level with respect to current realizations of inflation.

We note that this result is perhaps even more puzzling in light of **OU1**: If aggregate forecasts under-react to aggregate revisions, which are related to news about inflation, why would they over-react to current realizations of (i.e. news about) inflation?

3.4 D1: Disagreement

Turning our attention away from over/under-reaction for now, we verify key empirical facts about disagreement.

First, a simple visualization (Figure 2) reveals substantial differences in inflation expectations at any time.

[3] We are keenly aware that the set of coefficients is atypical. We are unable to verify if this result was obtained in other papers - Kohlhas and Walther (2018) did not run regressions on sub-samples, while Angeletos, Huo and Sastry (2020) only provided coefficients for the full sample and the post-1984 period.

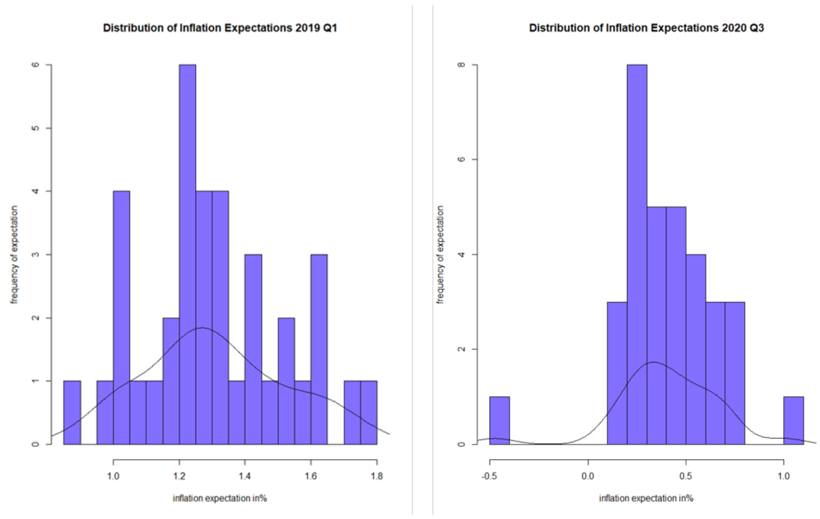

Figure 2: Distribution of inflation expectations at 2019 Q1 and 2020 Q3

Second, disagreement in inflation forecasts, which can be measured by the IQR (75th percentile - 25th percentile), varies substantially over time. For Figure 3, we use quarterly inflation forecasts, $\bar{\mathbb{E}}_t[\pi_{t-1,t,4,1}]$.

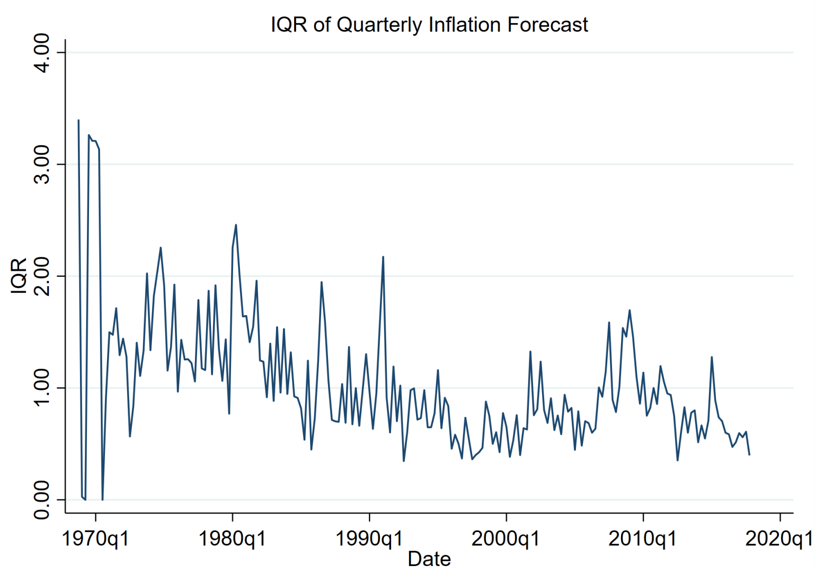

Figure 3: IQR of Quarterly Inflation Forecasts

Third, disagreement increases temporarily with both positive and negative shocks; this is seen in the Sticky-Information Model (Mankiw and Reis, 2002), which generates endogenous disagreement in expectations that is correlated with macroeconomic variables.

4 New Candidate Empirical Fact

4.1 OU4: Over-reaction and Under-reaction across Time

We saw both over-reaction and under-reaction in our established empirical facts **OU1**, **OU2** and **OU3**, which appear inconsistent. However, it is possible for both over- and under-reaction to affect expectations, but in different magnitudes across time. To identify how reaction varies across time, an intuitive concept would thus be an Impulse Response Function (IRF): How will inflation forecast errors vary over time with a 'cost-push' (or inflationary) shock?

4.1.1 SVAR

To compute an IRF, we first need to identify a cost-push shock. Our contribution here is to do so based on the methodology described by Angeletos, Collard and Dellas (2020), but with inflation forecasts as an additional variable.

Shocks are estimated from an 11-variable Structural Vector Autoregression (SVAR), containing: the inflation forecasts by SPF forecasters ($\bar{\mathbb{E}}_t[\pi_{t-1,t+3,1,1}]$), the unemployment rate, real GDP per capita, real investment per capita, real consumption per capita, hours worked per person, labour productivity in the non-farm business sector, level of utilisation-adjusted total factor productivity (TFP), labour share, inflation rate (GDP deflator) and the Federal Funds Rate.

The VAR was estimated with Bayesian methods using a Minnesota prior, due to its large size. The posterior distributions were obtained using Gibbs sampling with 50,000 draws. The reported 95% Highest Posterior Density Intervals (HPDI), or the Bayesian analog of frequentist 95% Confidence Intervals, were obtained by the approach described in Koop (2003). We use 2 lags, per standard Bayesian criteria.

The VAR takes the following linear form, by Wold Representation:

$$\boldsymbol{A}(L)\boldsymbol{Y_t} = \boldsymbol{u_t} \tag{8}$$

where $\boldsymbol{Y_t}$ is a 11×1 vector containing the variables listed above, $\boldsymbol{A}(L) = (\boldsymbol{I_n} + \boldsymbol{A_1}L + \boldsymbol{A_2}L^2)$ is the matrix polynomial in the lag operator L and $\boldsymbol{u_t}$ denotes the vector of VAR residuals with $E(\boldsymbol{u_t u_t'}) = \boldsymbol{\Sigma}$, where $\boldsymbol{\Sigma}$ is the positive definite covariance matrix of the VAR residuals.

We assume that the VAR residuals $\boldsymbol{u_t}$ are linear combinations of mutually independent "structural" shocks $\boldsymbol{\varepsilon_t}$:

$$\boldsymbol{u_t} = \boldsymbol{S\varepsilon_t} \tag{9}$$

where $\boldsymbol{S}$ is an invertible 11×11 matrix and $E(\boldsymbol{\varepsilon_t \varepsilon_t'}) = \boldsymbol{I}$.

The typical SVAR exercise in the literature employs exclusion or sign restrictions to identify the structural shock $\boldsymbol{\varepsilon_t}$. This implies imposing restrictions on the orthonormal matrix $\boldsymbol{Q}$, defined by the following equation:

$$\boldsymbol{S} = \boldsymbol{CQ} \tag{10}$$

where $\boldsymbol{C}$ is the Cholesky decomposition of $\boldsymbol{\Sigma}$, i.e. $\boldsymbol{\Sigma} = \boldsymbol{C}\boldsymbol{C}'$.

Here, however, a shock is identified by the requirement that it contains the maximal contribution to the volatility of a particular variable over a particular frequency band. This builds on the max-share approach (Uhlig, 2003; Barsky and Sims, 2011), by allowing the targeted variable and/or the targeted frequency band to vary, instead of committing to a specific choice. We consider unemployment and inflation as target variables; we also consider a targeted frequency band of $[2\pi/32, 2\pi/6]$, corresponding to frequencies between 6 and 32 quarters, typically associated with the business cycle.

It is noteworthy that this method shares a connection with that of Principal Component Analysis (PCA), whereby the shock ('factor') chosen is the eigenvector associated with the largest eigenvalue of a certain matrix; however, unlike PCA, it does not involve choosing any 'optimal' weights.

4.1.2 Shocks

Two shocks are identified:

The first shock, which contains the maximal contribution to the business cycle variation in unemployment, is called the 'Main Business Cycle' (MBC) shock. It has very similar IRFs and forecast error variance contributions to shocks identified by targeting many other macroeconomic variables, including hours worked, GDP, investment and consumption. As it has a small impact on inflation, it can be interpreted as a non- or mildly-inflationary 'demand' shock ε_t^D. We use this shock to check that the estimates from our SVAR are consistent with macroeconomic theory.

The second shock contains the maximal contribution to the business cycle variation in inflation and is nearly orthogonal to the MBC shock in the data. It can be interpreted as a 'cost-push' shock ε_t^S. We use this to generate IRFs of inflation actuals, inflation expectations and inflation forecast errors later on.

We adopt this methodology because a cost-push shock can be constructed to maximize the business-cycle variation in inflation. As seen in Table 4, the cost-push shock drives a substantial 71.4% of the business cycle variation in inflation - this is significantly higher than more standard technology (Gali, 1999), oil (Hamilton, 1996) and news (Barsky and Sims, 2011) shocks, each of which explains approximately 10-30% of the business cycle variation in inflation (Coibion and Gorodnichenko, 2012). Given that changes in inflation affect inflation expectations, it is thus likely that the cost-push shock will also drive a considerable component of the variation in inflation expectations - indeed, we see that the share of variation explained is 39.7%.

Table 4: Variance Contributions

Shock:	u	π	π^e
ε_t^D	70.6	9.7	12.6
	[57.4, 78.5]	[2.6, 18.2]	[4.7, 28.4]
ε_t^S	10.5	71.4	39.7
	[6.6, 13.4]	[63.6, 83.8]	[29.5, 49.3]
Note:		95% Reported HPDI	

4.1.3 IRFs using SVAR

We first look at the IRFs of the MBC shock on macroeconomic variables as a check.[4] They largely fall in line with what macroeconomic theory predicts: a positive MBC shock induces an initial rise in inflation expectations, a fall in unemployment, an increase in output, hours worked, and investment, as well as a rise in nominal interest rates. Other variables such as consumption, TFP, labour productivity, labour share, and inflation, experience ambiguous or negligible effects.

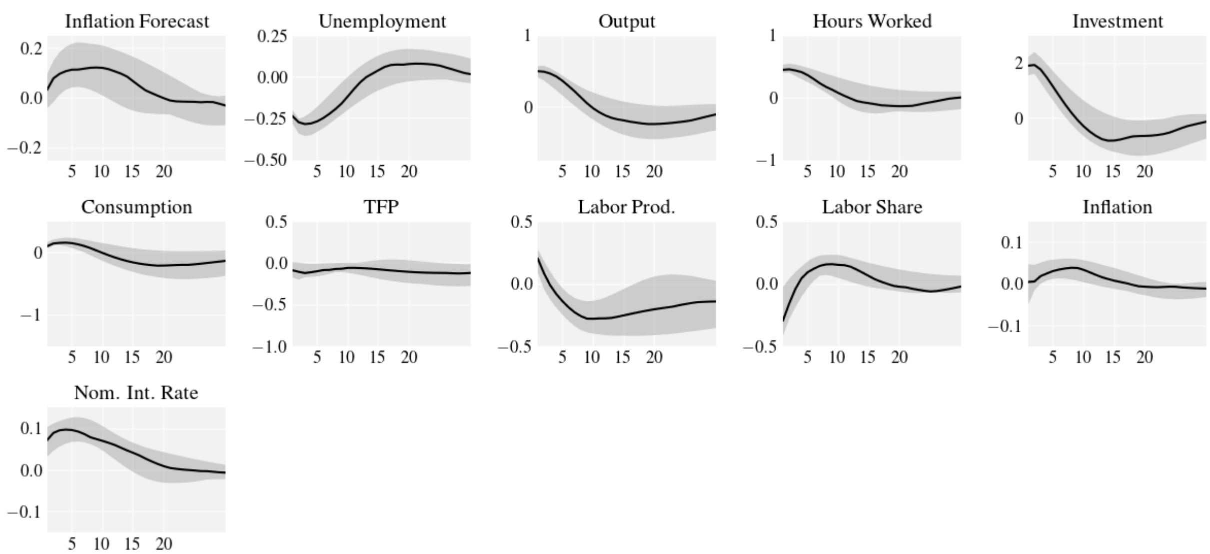

Figure 4: IRF of Macroeconomic Variables with MBC shock (95% HPDI)

We next look at the IRFs of the 'cost-push' shock on macroeconomic variables. The positive cost-push shock is estimated to raise inflation, inflation expectations and unemployment, as well as lower output, investment, consumption and TFP, which is consistent with theory.

[4]For all IRF plots in this paper, the x-axis denotes quarters from the shock (starting at 0), and the y-axis denotes the variable of interest. Note that the initial IRF coefficient for unemployment is not necessarily 1, because the MBC shock is not an unemployment shock, but a shock with the maximal contribution to the volatility of business cycle unemployment; likewise for the cost-push shock.

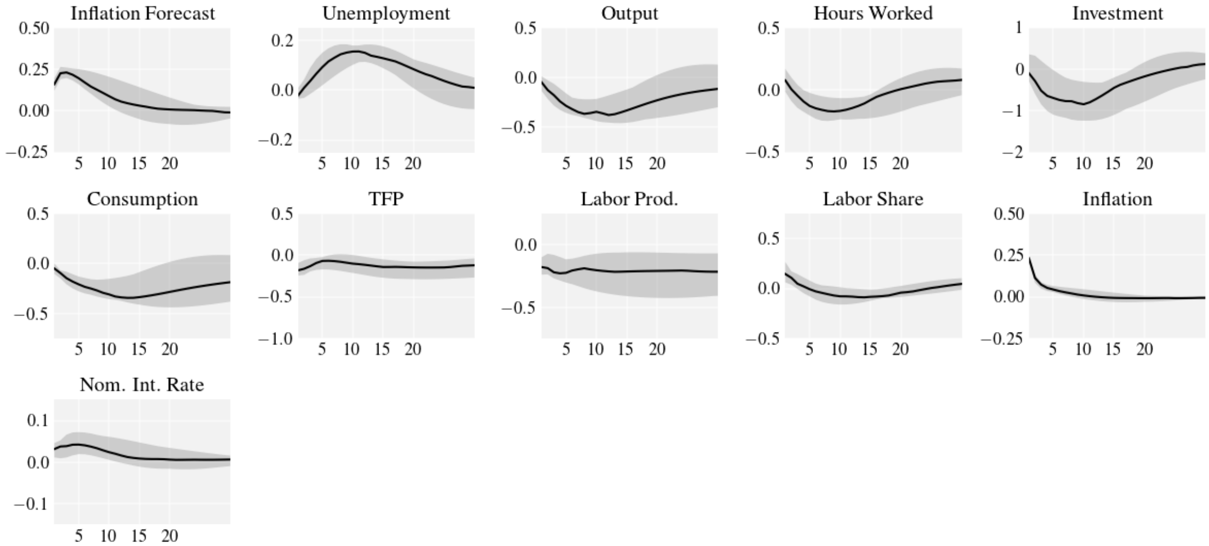

Figure 5: IRFs of Macroeconomic Variables with Cost-Push shock (95% HPDI)

We can obtain the IRF of inflation forecast errors by taking the difference between the IRF of inflation and inflation forecasts (Angeletos, Huo and Sastry, 2020). Scaling the IRF appropriately to achieve an initial coefficient of 1 for inflation, and comparing the IRFs of inflation, inflation forecasts, and inflation forecast errors gives us the following figure:

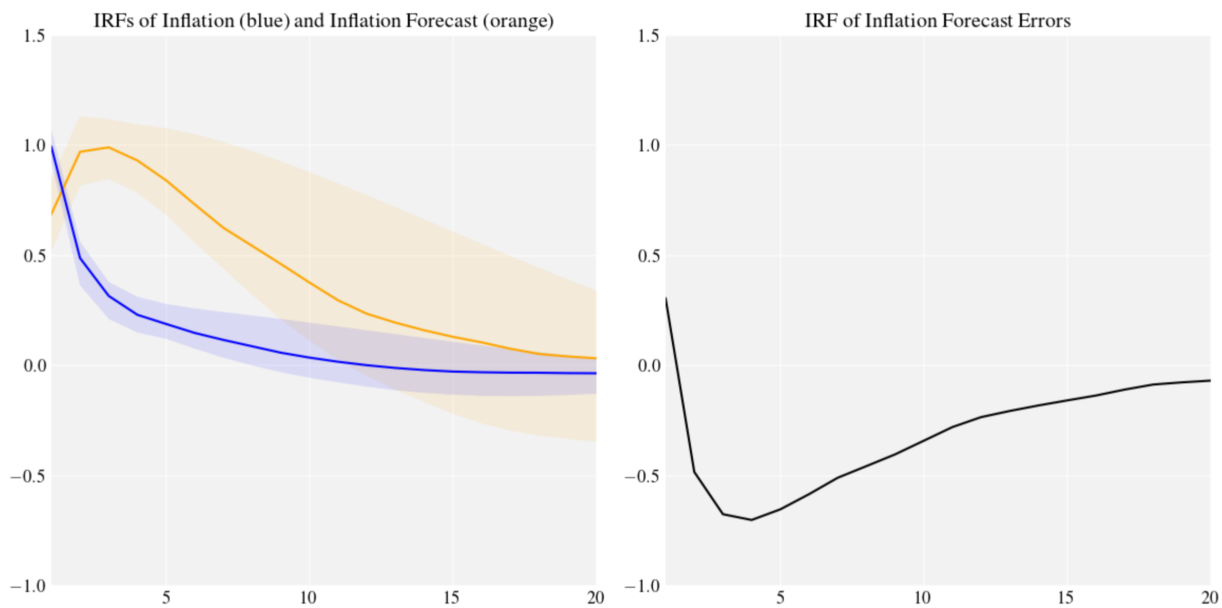

Figure 6: IRFs of Inflation Actuals, Inflation Forecasts and Inflation Forecast Errors, under the modified 11-variable ACD SVAR method (95% HPDI)

Our key object of interest is the IRF of inflation forecast errors. We see possible evidence of an interesting finding: that the IRF coefficients are initially positive, but subsequently turn negative.

We note that the sign switch occurs just 1 quarter after the shock here, which would imply that agents only under-react initially, and over-react from the next quarter onwards. We also note that along the horizon, the HPDI widens and the associated SEs increase, which could call into question the strength of the result.

4.1.4 IRFs using LPM

Angeletos, Huo and Sastry (2020) chose instead to run the original 10-variable ACD SVAR without inflation forecasts, derive a similar 'cost-push' shock, and then estimate the IRFs of interest using ARMA-IV and the Local Projection Method (LPM). As ARMA-IV uses lags as instruments, and we do not believe sequential exogeneity holds in our case[5], we choose to verify our results using the LPM instead. The LPM, developed by Jordà (2005), runs the following set of regressions for the horizons $h \in \{0, 1, ..., H\}$:

$$z_{t+h} = \alpha_h + \beta_h \cdot \varepsilon_t^S + \gamma' W_t + u_{t+h} \tag{11}$$

where $z_t \in \{\pi_{t-1,t+3,1,1}, \bar{\mathbb{E}}_t[\pi_{t-1,t+3,1,1}]\}$ (i.e. inflation actuals or inflation forecasts), ε_t^S is the 'cost-push' shock, W_t is a vector of controls, and γ is a vector of coefficients on the controls. We set H=20 quarters here. Controls include 4 lags of each of the 10 macroeconomic variables that entered the original SVAR.

Implementing the LPM gives us the following IRFs for inflation, inflation forecasts, and inflation forecast errors:

[5] An instrument is valid if it is not correlated with errors (u_t), in a regression of inflation/inflation expectations (z_t) on the cost-push shock (ε_t^S). Here, the instrument proposed is the lagged shock, ε_{t-k}^S, with $k \in \{3, ..., 11\}$. So we need, for instance, Corr(ε_{t-3}^S, u_t)=0. This is satisfied under sequential exogeneity, in which inflation/inflation expectations today don't predict errors tomorrow.

However, we contend otherwise: A cost-push shock can affect inflation expectations for many periods ahead, as seen in Enders, Hünnekes and Müller (2019), so Corr(ε_{t-3}^S, $\bar{\mathbb{E}}_{t-3}[\pi_{t-1,t+3,1,1}]$)$\neq$0. Firms may then respond to changing inflation expectations by changing their investment patterns, which would then affect future growth and inflation, as seen in Coibion, Gorodnichenko and Ropele (2020), so Corr($\bar{\mathbb{E}}_{t-3}[\pi_{t-1,t+3,1,1}]$, $\pi_{t-1,t+3,1,1}$)$\neq$0. And finally, inflation feeds back into expectations, so Corr($\pi_{t-1,t+3,1,1}$, $\bar{\mathbb{E}}_t[\pi_{t-1,t+3,1,1}]$)$\neq$0. In other words, Corr($\varepsilon_{t-3}^S$, z_t)$\neq$0. Since we know that Corr(z_t, u_t) $\neq 0$, we have that Corr(ε_{t-3}^S, u_t)$\neq$0, and lagged shocks are thus invalid instruments for both inflation and inflation expectations.

As Blundell and Bond (1998) pointed out, there are many issues with lags as instruments, especially that of weak IV (i.e. instruments are not relevant). In fact, we see this in the results by Angeletos, Huo and Sastry (2020): the reported first stage F statistics for the ARMA-IV estimates were 2.1 for inflation outcomes and 2.4 for inflation forecasts, significantly below the heuristic of 10 proposed by Stock and Yogo (2005). This concern about weak IV was also the authors' motivation for implementing alternative methods, such as the LPM.

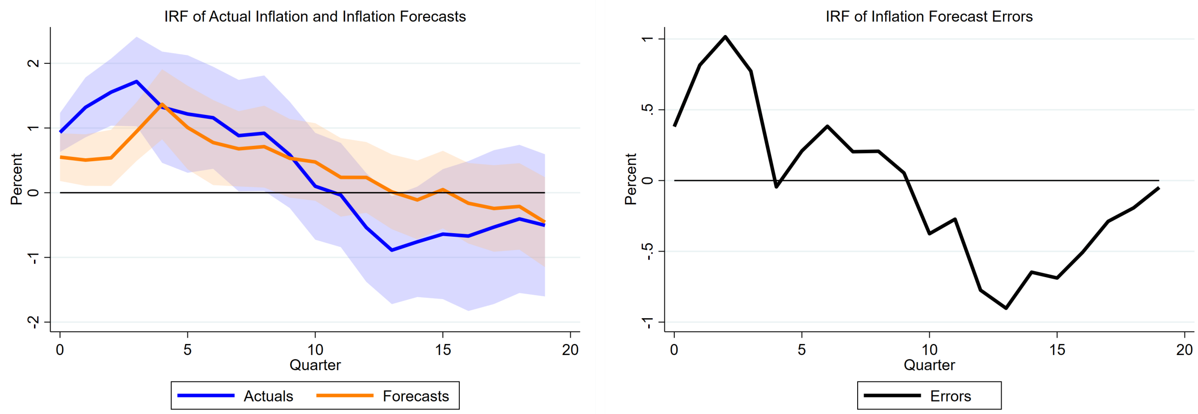

Figure 7: IRFs of Inflation Actuals, Inflation Forecasts and Inflation Forecast Errors, under the Local Projection Method. (95% CI, RSEs)

We see a clearer pattern of sign reversal in the IRF of inflation forecast errors, and our result here is significantly closer to that obtained by Angeletos, Huo and Sastry (2020). While this builds a case for **OU4**, we view this candidate empirical fact as suggestive rather than definitive, in light of the high uncertainty at longer horizons and the less striking results obtained from the SVAR.

5 A Possible Unified Explanation

5.1 Model

Finally, we would like to identify a unified explanation for deviations from FIRE as observed in **OU1**, **OU2**, **OU3**, **OU4** and **D1**. We do this using a simplified model (Angeletos, Huo and Sastry, 2020; Reis, 2020) and leveraging it to isolate plausible mechanisms. At this level of generality, we collapse the notation for quarterly inflation to $\pi_{t+s} = \pi_{t+s-1,t+s,k,l}$ for a consistent $[k, l]$.

We start off with a state of the world where π_t follows an AR(1) process with Gaussian errors: $\pi_t = \rho\pi_{t-1} + \epsilon_t$, where $\rho < 1$ represents persistence of the process (past innovations) and $\epsilon_t \sim N(0, 1)$ is a Gaussian innovation.

In a FIRE world, all individuals are aware of actual inflation, and thus the signal observed by each individual i in period t is simply $s_{i,t} = \pi_t$.

We then consider deviations from FIRE, firstly with respect to full information. We introduce a channel for information frictions, in the form of *dispersed, noisy information*: an agent's observation of π_t is contaminated with idiosyncratic noise. Each individual i thus observes a signal in period t, given by:

$$s_{i,t} = \pi_t + \frac{u_{i,t}}{\sqrt{\tau}} \tag{12}$$

where τ measures precision and $u_{i,t} \sim N(0, 1)$ is i.i.d. Gaussian noise. In other words, individuals receive different signals about inflation, and are thus heterogeneous in how well-informed they are[6]. In light of noisy signals, we assume individuals form inflation expectations via a Kalman filter:

$$\mathbb{E}_t(\pi_t)_i = g s_{i,t} + (1-g)\,\mathbb{E}_{t-1}(\pi_t)_i \tag{13}$$

where $g \in [0,1]$ is the Kalman gain, an increasing and continuous function of τ, with $g \to 0$ as $\tau \to 0$ and $g \to 1$ as $\tau \to \infty$. As the precision of the signal increases, individuals view the signal as being more reliable, and place a greater weight on it when forming inflation expectations. If precision is known to be low, then individuals fall back more onto their previous forecast instead. If $g < 1$, then individuals do not update sufficiently to new information, and thus under-react to changes in inflation.

We also introduce a channel for *sticky information*: suppose a fraction $(1-\theta)$ of agents have limited information, and form expectations according to $\mathbb{E}_t(\pi_t)_i = g s_{i,t} + (1-g)\,\mathbb{E}_{t-1}(\pi_t)_i$. However, a fraction θ have full information, so their forecast is instead $\mathbb{E}_t(\pi_t)_i = \pi_t = \rho\pi_{t-1}$.

Following which, we consider deviations from rational expectations, by introducing channels for perceived persistence and precision to be different from actuals.

This means that agents perceive the true process to be:

$$\pi_t = \hat{\rho}\pi_{t-1} + \epsilon_t \tag{14}$$

where $\hat{\rho}$ is the **perceived** persistence. If $\hat{\rho} > \rho$, agents overestimate the persistence of the process, and thus the effect of an innovation today on inflation tomorrow; in other words, *over-extrapolation*. (This would be clearer by shifting forward the true process by 1 period: $\pi_{t+1} = \rho\pi_t + \epsilon_{t+1}$.) The opposite case, $\hat{\rho} < \rho$ represents *under-extrapolation*.

Agents also perceive the process of the signal to be:

$$s_{i,t} = \pi_t + \frac{u_{i,t}}{\sqrt{\hat{\tau}}} \tag{15}$$

where $\hat{\tau}$ is the **perceived** precision. If $\hat{\tau} > \tau$, agents overestimate the precision of the signal process, and are thus *overconfident* about the signal. The opposite case, $\hat{\tau} < \tau$ represents *under-confidence*. Estimates by Angeletos, Huo and Sastry (2020) did not find a significant role for over-confidence given over-extrapolation as an explanatory factor, and we thus leave it out for parsimony.

5.2 Reconciling Model with Empirical Facts on Over/Under-reaction

This model gives us the ability to isolate mechanisms that support empirical facts **OU1**, **OU2**, **OU3** and **OU4**.

We saw that from **OU2**, a regression of individual errors on individual forecast revisions, $\beta_{BGMS} < 0$. However, we note that under rational expectations, an individual's forecast error is unpredictable by his own past information, regardless of information frictions. This is because rational agents

[6]From a broader perspective, this can also be interpreted as Rational Inattention.

with limited information would not systematically make errors in light of inflation news. It is thus indicative that deviations from RE, in the form of over-extrapolation ($\hat{\rho} > \rho$), is required to explain our result. If agents overestimate the effect of any given innovation today on future outcomes, then they systematically over-react. This also implies that models predicting under-extrapolation at the individual level, such as level-K thinking, are inconsistent with the empirical fact.

We also saw that from **OU1**, a regression of aggregate errors on aggregate forecast revisions, $\beta_{CG} > 0$. With only over-extrapolation, we should expect over-reaction at the aggregate level, which is counterfactual. Thus, we see a role for information frictions here: with the effect of $u_{i,t}$ working through $\tau > 0$ and the Kalman gain $g < 1$, dispersed noisy information can lead to under-reaction at the aggregate level.

The two mechanisms of information frictions coupled with over-extrapolation also corroborate with **OU4**, which revealed that the IRF of average errors is initially positive, but subsequently becomes negative.

The IRF of the average forecast error can be defined as a plot of $\{\zeta_k\}_{k=0}^{\infty}$, where ζ_k is the k^{th} coefficient in the infinite moving-average representation of the average forecast error:[7]

$$\zeta_k(\rho, \hat{\rho}, \tau, \hat{\tau}) = \frac{\partial Error_{t+k}}{\partial \epsilon_t} = \frac{\partial(\pi_{t+k-1,t+k+3} - \bar{\mathbb{E}}_{t+k}[\pi_{t+k-1,t+k+3}])}{\partial \epsilon_t} \tag{16}$$

Angeletos, Huo and Sastry (2020) mathematically showed the key result of under-reaction and over-reaction with the following proposition: If $\hat{\rho} > \rho$ and $\hat{\tau}$ is small enough relative to $(\hat{\rho} - \rho)$, or agents over-extrapolate and information frictions are large enough (learning is slow enough), then $\zeta_k > 0$ for $0 \leq k < k_{IRF}$ (i.e. initial under-reaction) and $\zeta_k < 0$ for $k > k_{IRF}$ (i.e. subsequent over-reaction), for some $k_{IRF} \in (0, \infty)$.

This sign-switch thus provides evidence of both under-reaction (as predicted by information frictions) and over-reaction (as predicted by over-extrapolation). Shortly after a shock occurs, informational frictions imply that forecasts under-react. However, as time passes and individuals learn, over-extrapolation takes over, implying that forecasts eventually over-react. This behaviour is also consistent with an intuitive story: First, agents form initial beliefs given available information. Then, news arrives and challenge initial beliefs - however, there is often inherent psychological inertia against changing beliefs. This manifests as doubt regarding the informativeness of new data, leading to initial under-reaction. But as the shock subsides and the truth becomes apparent, agents overcome the psychological barrier and update their beliefs, often in a significant way, thus leading to over-reaction.

We are left with **OU3**, a regression of aggregate errors on current realizations, where $\beta_{KW} < 0$. First, we can derive an analytical expression for β_{CG} and β_{KW} in terms of the IRF of average forecast errors, the IRF of average forecast revisions, and the IRF of Gaussian innovations, by considering the infinite MA representations of forecast errors, forecast revisions, and actual outcomes. Angeletos, Huo and Sastry (2020) then show that β_{CG}, containing the dot-product of IRFs of forecast errors and revisions, has more positive terms than β_{KW}, containing the the dot-product of the IRFs of

[7] Note that the IRF is an object containing a mixture of present and future quarterly inflation, due to the inflation forecast horizon being h=3 (i.e. forecast for 4-period inflation). This is a significant complication that is largely glossed over in theoretical models, due to a preference for simpler notation such as π_{t+k} and π_{t+k-1}.

forecast errors and outcomes. As such, β_{CG} can be positive, while β_{KW} can be negative under information frictions and over-extrapolation. It can also be interpreted the following way: β_{CG} puts more weight on the early positive portion of the IRF of forecast errors, while β_{KW} puts more weight on the later negative portion.

This model, together with information frictions and over-extrapolation, could thus explain all 4 empirical facts about over- and under-reaction. However, we highlight the difficulty in interpreting and drawing meaning from OLS results, in particular β_{CG} and β_{KW}, as they collapse time-varying coefficients into a less informative point estimate, and even do so in different ways.

5.3 Reconciling Model with Empirical Facts on Disagreement

Reis (2020) mathematically showed that the law of motion for disagreement in this model is a function of the shock ϵ_t:

$$(1-\hat{\lambda})^2 V_t = \theta(1-\theta)\hat{\lambda}^2\epsilon_t^2 + (1-\theta)(\hat{\rho}-\hat{\lambda})^2 \tag{17}$$

where V_t is disagreement measured by the (cross-sectional) variance of inflation expectations at time t, and $\hat{\lambda} \in (0, \hat{\rho})$ is the root of the quadratic: $\hat{\lambda} + 1/\hat{\lambda} = \hat{\rho} + (1+\tau)/\hat{\rho}$. The model thus captures **D1**, as the ϵ_t^2 term ensures that disagreement changes over time and increases temporarily with both positive and negative shocks.

This completes the model, which provides a unifying explanation for all empirical facts: **OU1**, **OU2**, **OU3**, **OU4** and **D1**.

6 Conclusion

Deviations from FIRE are evident, but a unifying explanation remains elusive. We have presented a model of expectations grounded by Angeletos, Huo and Sastry (2020) and Reis (2020) to capture empirical facts about both over/under-reaction and disagreement.

There are important caveats to be aware of: (1) the sign reversal finding from **OU4** is not necessarily conclusive due to high SEs at long horizons; (2) empirical 'facts' may not capture important aspects of reality; and that (3) empirical 'facts' are subject to the Lucas critique: if individuals change their expectation formation processes with significant events, and events happen with either sufficient speed or surprise, it is unclear what 'facts' can actually consistently underpin theoretical frameworks over time.[8]

Nonetheless, the IRF provides a significantly stronger and (perhaps surprisingly) more intuitive way to understand inflation expectations data than standard OLS regressions, because it bet-

[8] There are models dealing with deviations from rational expectations, which are not subject to the Lucas critique. For instance, diagnostic expectations generate over-reaction via the representativeness heuristic (Bordalo, Gennaioli, Ma, and Shleifer, 2020), while memory-based models generate over-reaction via recent information being more 'available' in mind (Afrouzi, Kwon and Ma, 2020; da Silveira, Sung and Woodford, 2020). These models do not depend on the past mechanically, and could thus mitigate, if not resolve stationarity issues.

ter reflects the complexity behind expectation formation processes. By allowing us to see how over/under-reaction changes over time, we gain insight on the potential for multiple competing mechanisms at work, and to isolate mechanisms that corroborate with the pattern of IRFs. The proposed parsimonious model of expectations is also a relatively tractable one, that can serve as the starting point for more nuanced and context-specific considerations.

Finally, further direction of work could lie in examining if the proposed theories laid out in this paper apply empirically in financial markets; specifically, (1) whether implied inflation rates, as measured by inflation-linked swaps and bonds, tally with predictions in terms of sign and magnitude, (2) whether prices of general assets such as equities, bonds and derivatives respond similarly to that of inflation-linked securities, and (3) whether the price of dividend-paying assets such as equities respond to a dividend shock in an analogous way.

References

Afrouzi, Hassan, Spencer Kwon, and Yueran Ma. "A Model of Costly Recall". *Working Paper, Columbia University* (2020).

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas. "Business-Cycle Anatomy". *American Economic Review* 110, no. 10 (2020): 3030–70.

Angeletos, George-Marios, Zhen Huo, and Karthik A Sastry. "Imperfect Macroeconomic Expectations: Evidence and Theory". Forthcoming, *NBER Macroeconomics Annual* (2020).

Angeletos, George-Marios, and Chen Lian. "Incomplete Information in Macroeconomics: Accommodating Frictions in Coordination". *Handbook of Macroeconomics* 2 (2016): 1065–1240.

Barsky, Robert B, and Eric R Sims. "News Shocks and Business Cycles". *Journal of Monetary Economics* 58, no. 3 (2011): 273–289.

Blundell, Richard, and Stephen Bond. "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models". *Journal of Econometrics* 87, no. 1 (1998): 115–143.

Bordalo, Pedro, et al. "Overreaction in Macroeconomic Expectations". *American Economic Review* 110, no. 9 (2020): 2748–82.

Broer, Tobias, and Alexandre Kohlhas. "Forecaster (Mis-) Behavior" (2018).

Clarida, Richard, Jordi Gali, and Mark Gertler. "The Science of Monetary Policy: A New Keynesian Perspective". *Journal of Economic Literature* 37, no. 4 (1999): 1661–1707.

— . "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory". *The Quarterly Journal of Economics* 115, no. 1 (2000): 147–180.

Coibion, Olivier, and Yuriy Gorodnichenko. "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy* 120, no. 1 (2012): 116–159.

— . "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts". *American Economic Review* 105, no. 8 (2015): 2644–78.

Coibion, Olivier, Yuriy Gorodnichenko, and Tiziano Ropele. "Inflation Expectations and Firm Decisions: New Causal Evidence". *The Quarterly Journal of Economics* 135, no. 1 (2020): 165–219.

Da Silveira, Rava Azeredo, Yeji Sung, and Michael Woodford. "Optimally Imprecise Memory and Biased Forecasts". *NBER Working Paper*, no. w28075 (2020).

Enders, Zeno, Franziska Hünnekes, and Gernot J. Müller. "Monetary Policy Announcements and Expectations: Evidence from German Firms". *Journal of Monetary Economics* 108 (2019): 45–63.

Gabaix, Xavier. "A Behavioral New Keynesian Model". *American Economic Review* 110, no. 8 (2020): 2271–2327.

Gali, Jordi. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review* 89, no. 1 (1999): 249–271.

García-Schmidt, Mariana, and Michael Woodford. "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis". *American Economic Review* 109, no. 1 (2019): 86–120.

Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer. "Expectations and Investment". *NBER Macroeconomics Annual* 30, no. 1 (2016): 379–431.

Greenwood, Robin, and Andrei Shleifer. "Expectations of Returns and Expected Returns". *The Review of Financial Studies* 27, no. 3 (2014): 714–746.

Hamilton, James D. "This is What Happened to the Oil Price-Macroeconomy Relationship". *Journal of Monetary Economics* 38, no. 2 (1996): 215–220.

Jordà, Òscar. "Estimation and Inference of Impulse Responses by Local Projections". *American Economic Review* 95, no. 1 (2005): 161–182.

Kiley, Michael T. "A Quantitative Comparison of Sticky-Price and Sticky-Information Models of Price Setting". *Journal of Money, Credit and Banking* 39 (2007): 101–125.

Kohlhas, Alexandre, and Ansgar Walther. "Asymmetric Attention". *Available at SSRN 3161628* (2018).

Koop, Gray. *Bayesian Econometrics.* Hoboken, NJ: John Wiley & Sons, 2003.

Maćkowiak, Bartosz, Filip Matějka, and Mirko Wiederholt. "Survey: Rational Inattention, a Disciplined Behavioral Model" (2018).

Mankiw, N Gregory, and Ricardo Reis. "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve". *The Quarterly Journal of Economics* 117, no. 4 (2002): 1295–1328.

Mankiw, N Gregory, Ricardo Reis, and Justin Wolfers. "Disagreement about Inflation Expectations". *NBER Macroeconomics Annual* 18 (2003): 209–248.

Reis, Ricardo. "Imperfect Macroeconomic Expectations: Yes, But, We Disagree" (2020).

Sims, Christopher A. "Implications of Rational Inattention". *Journal of Monetary Economics* 50, no. 3 (2003): 665–690.

Stock, James H, and Mark W Watson. "Forecasting Output and Inflation: The Role of Asset Prices". *Journal of Economic Literature* 41, no. 3 (2003): 788–829.

Stock, James H, and Motohiro Yogo. "Testing for Weak Instruments in Linear IV Regression". *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg* 80, no. 4.2 (2005): 1.

Uhlig, Harald. "What moves real GNP?" (2003).

Woodford, Michael. "Imperfect Common Knowledge and the Effects of Monetary Policy". *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps* (2003).

WIC and Maternal Labour Market Effects in the US *

Mina Burns Ricky Wang Marco Brolli Shawn Teo

August, 2021

Abstract

The Special Supplemental Nutrition Program for Women, Infants, and Children (WIC) in the United States provides food-benefits to low income women and children; children are disqualified for WIC benefits when they turn 5. To study the effect of WIC on the maternal labour market and household food insecurity, we use age of the youngest child as a proxy for WIC participation to employ a fuzzy regression discontinuity design. Using data from the Current Population Survey (CPS), we study maternal employment and hourly wage as well as whether a household is very food insecure. We find evidence that participation in WIC does not appear to have an effect on maternal employment, hourly wage and whether a household is very food insecure. Exit from WIC does appear to a significant effect on whether children in a household are food insecure. Furthermore, we propose potential explanations for these findings.

*Department of Economics, London School of Economics and Political Science, London, UK. Corresponding email address: c.e.burns@lse.ac.uk. This research paper was a team effort and we are very grateful for the support from the LSE Economics Society. Generous thanks to Professor Lucie Schmidt at Williams College for answering our (many!) questions.

1 Introduction

The Covid-19 pandemic and its damaging effects on the economy have exacerbated the pre-existing food-security issues in the United States. According to the U.S. Department of Agriculture, the pandemic has caused over 30 million adults and, due to in part to school closures, 12 million children to live in a food insecure environment. The lack of access to food of decent quality has moreover worsened the country's racial disparity. In fact, Black and Latino people are more than twice as likely to not have enough to eat compared to white people (USDA, Release No. 0037.21). To change the course of events, the Biden administration has pledged, as part of the American Rescue Plan, to increase the Supplement Nutrition Assistance Program (SNAP) benefits by 15%. This paper explores a similar program, the WIC program, which targets a very specific subset of the population, women and children. While, the link between the WIC and its beneficial consequences on food security is already supported by strong evidence (Herman et al., 2004 and Kreider et al., 2016), our intention is to understand whether such programs aimed mainly at food and poverty alleviation could have a more long-lasting economic impact on the participants. In particular, we seek to understand whether the alleviation of food insecurity could help mothers to enter the labour market.

2 Literature Review

The Women, Infants and Children (WIC) program aims to provide high quality food to low-income pregnant women, mothers, and children up to 5 years of age, focusing on improving and safeguarding their health. The program, in most states, works by handing participants paper vouchers or electronic benefit cards redeemable at authorized retail stores for a certain set of nutrient-rich supplemental foods. WIC serves 53% of all infants and 25% of all pregnant women in the United States: this accounts for over 9 million mothers and children each month (Center for Hunger Free Communities, report, 2012)

With a cost of over 5.3 billion dollars in the 2018 fiscal year, WIC is one of the most extended food welfare programs of the United States. In fact, in April 2018, there were over 7.8 million people certified to receive benefits from (WIC). While the number of participants has steadily increased from 1992 through 2010, the number of participants has significantly decreased since 2010. Participation declined by 12.0 percent from 2010 to 2016 and an additional 11.1 percent from 2016 to 2018 (USDA, 2018).

Most WIC participants come from households with income at or below 185 percent of the Federal Poverty Guidelines issued each year by the U.S. Department of Health and Human Services. These income guidelines vary by household size and residency and are used to determine eligibility for certain programs including WIC. The average annualized income of WIC households in 2018 was $19,355. Across participant categories, household income was highest for breastfeeding women ($20,343) and lowest for postpartum women ($16,631) in 2018. Across racial and ethnic groups, average household income was highest for Asian participants ($23,999) and lowest for Black participants ($15,033)(WIC report, 2018).

Most previous studies have focused on WIC's positive effects on maternal health and infant health as well as on quantifying the savings for the healthcare system. The literature has suggested that every $1.00 spent on WIC saves of up to $4.21 for the welfare system thanks to WIC's 25% reduction in the risk of preterm birth and 44% reduction in low birth weight babies (National WIC Association). Research concludes that WIC improved several other key birth outcomes. The average first year medical costs for a baby born without complications are $4,551 compared to the average first year medical costs of $40,003 for a preterm or low birth weight baby.

Understanding what prevents access to a welfare program is crucial for its effectiveness. Often times, excessive red-tape or inefficiency could exacerbate the so-called welfare stigma (see Moffitt, 1983 and Coate et al., 1992).This, in turn could lead to a worsening of the psychological situation of the recipient (Pak, 2020), further damaging the recipient's economic situation. For this reason, much of the previous literature on the topic looked at the barriers to the WIC program. In particular, linguistic challenges, lack of cultural competencies, issues related to the clientâstaff interaction and that insufficient time, materials, and staffing could all prevent women the access to WIC (Seth et al. 2015). Research has also shown that structural barriers such as transportation, difficulties in obtaining appointments, or the inability to take time off work can result in a significant reduction in WIC participation for WIC-eligible women. WIC-eligible women reporting unplanned pregnancies and fewer supports from their familiar or social network, also tend to participate less in WIC. (C. Liu & H. Liu, 2016).

In order to get a proper grasp of how the WIC program should act on the labour market and on the labour supply of the women involved, we should start by analysing the effects of other welfare policies on the labour supply curve. From a theoretical standpoint, it could be hypothesised that better nutrition alongside an improved metal health arising from a partial relief of the burden from procuring quality food might improve the productivity of the mother, allowing her to supply more hours of work. It could also be argued that, by reducing the cost of remaining unemployed, the welfare benefit could reduce the incentive to work, with the opposite effect on the labour supply.

The empirical evidence offered mixed results, varying depending on the type of welfare support provided. On one side, we have the working families' Tax Credit, introduced in the UK to help reduce child poverty both by attracting parents from previously jobless households into the labour market and by distributing additional financial support to those already working but living in families with a low income. Its impact has been positive as it increased labour supply of lone mothers by around 5.1 percentage points, although slightly reduced labour supply of mothers in couples by 0.6 percentage points and increased the labour supply of fathers in couples by 0.8 percentage points. (Brewer et al., 2006). Similarly, another way to facilitate the entry of single women into the workforce has been the EITC (the Earned Income Tax Credit) which consisted of cash-benefit entitlement granting a wage higher than the actual one for the low-income earner. Meyer and Rosenbaum (2001) observed the EITC has been successful increasing the labour force participation of unmarried women with children by 2.4% and the effect has even been larger for women with younger kids. The greater the extent of public assistance through programs such as Supplemental Security Income and Food Stamps has been found to increase female labor market participation (Cebula and Coombs, 2008). On the other hand, there are examples where welfare benefits could instead decrease the labour supply of participants. For example, generous disability benefits reduce the labour supply of men nearing retirement age (Gruber, 2000). Similarly, a study on the effect of food stamps on the labour supply of immigrant concludes that access to the program reduces the employment rates of single women by about 6%, whereas married men continue to work but reduce their hours of work by 5% (East, 2018).

The Food Stamp Program, food insecurity, wage, and labor force participation have heavily interconnected relations. Huffman and Jensen (2008) found that labor force participation does not appear to have an affect on food insecurity, although it reduces participation in the Food Stamp Program. They also found that a higher wage decreases participation in Food Stamps and increase labour force participation. In addition, Food Stamp participation fails to have a significant relationship with food insecurity. These complex relationships will be further explored in this paper.

3 Data

The primary dataset we use is the IPUMS Current Population Survey (IPUMS-CPS) by Flood et al. (2020). This consists of harmonised microdata composed of individual records containing information collected on persons and households from the US Census Bureauâs monthly Current Population Survey (CPS), which is a rotating panel design with more than 65,000 households. Each household is rotated in for four consecutive months, rotated out for the next eight consecutive months, and then rotated in again for four consecutive months, meaning that households are in the CPS for a total of eight months. Beyond the core CPS, we also use the Food Security Supplement which further collects information including food spending, food security, and participation in various food assistance programs, including WIC. This Supplement is included as part of the CPS for households polled in December each year from 2001. In addition, we use GDP by industry data from the US Bureau of Economic Analysis (BEA) to determine which industries are growing.

Our data is from the years 2001-2019 and we also recode other variables of interest for ease of use within the regressions. Further, to avoid complications with the difference between infant WIC package compared to the child WIC package, we focus on children aged 1–8.

4 Identification Strategy

The WIC program has four eligibility requirements, namely:

1. Categorical - one of the following:
   (a) Women - pregnant, postpartum, or breastfeeding
   (b) Infants up to first birthday
   (c) Children up to fifth birthday
2. Residential - live in the state to which they are applying
3. Income - at least one of the following:
   (a) Below 185% federal poverty guidelines
   (b) SNAP, Medicaid, or TANF recipient
   (c) Recipient of other state program
4. Nutrition Risk - determined by health professional

We exploit the first requirement of a child below the age of five using a parametric regression discontinuity design (RDD), using the youngest childâs age as the running variable. However, since compliance is likely imperfect - not every person eligible for WIC will enroll in the program - a fuzzy RDD design is necessary. One concern with using this design is that the CPS data does not specify how someone qualifies for WIC. In other words, we do not know whether the WIC food package is received by a pregnant, postpartum, or breastfeeding woman or an infant or a child. In addition, the CPS data does not include data on whether a woman in the house is pregnant, postpartum, or breastfeeding. Assuming that all other characteristics are distributed continuously across the cutoff, those below and above the cutoff should be almost identical. Therefore, the difference in outcomes between the two groups would be solely due to enrolment in WIC.

Let D_i be assigned according to the cutoff:

$$D_i = \begin{cases} 1 \text{ if } A_i < 5 \\ 0 \text{ if } A_i \geq 5 \end{cases} \tag{1}$$

where A_i is the running variable, age of child. Let WIC_i be the treatment dummy, equal to 1 for those who receive WIC and 0 for those who do not. Let b be the bandwidth, so that only children with $5 - b \leq A_i \leq 5 + (b - 1)$ are considered.

Consider our first aim of examining the effect of WIC on the chance of a mother entering a high-growth industry. We use an instrumental variables estimation approach, with the first and second stage local linear regressions as follows:

$$WIC_i = \alpha_0 + \alpha_1 D_i + \alpha_2 f(A_i - 5) + \alpha_3 G_i + e_i \tag{2}$$

$$Y_i = \beta_0 + \beta_1 W\hat{I}C_i + \beta_2 f(A_i - 5) + \beta_3 G_i + \varepsilon_i \tag{3}$$

In the regressions above, G_i indicates various controls, as indicated below in Table 2. These regressions are then repeated for each of our outcomes of interest, replacing the outcome variable in each case.

To ensure the robustness of our results, we also include linear and non-linear trends, as indicated by $f(A_i - 5)$ in equations 2 and 3. In order to ensure that our sample reflects individuals that meet the WIC income eligibility requirement, we restrict our sample to households below the 185% poverty threshold. We also restrict our samples to households where the child is the youngest in the household in order to households that would experience discrete changes in WIC participation.

Another major concern, particularly for maternal labour market effects, is the effect of school attendance. While some children may enroll in school programs before the age of 5, school becomes mandatory for 5 year old children in many states. School meal programs further complicate school attendance. Many children exit the WIC program to enter school and enroll in the lunch or breakfast program. Varying by state, many schools provide free or reduced-cost lunches or breakfasts to its students. These meals vary by nutrition content and availability according to state and could possibly diminish the effect of WIC. We therefore further restrict our sample to children that don't attend school.

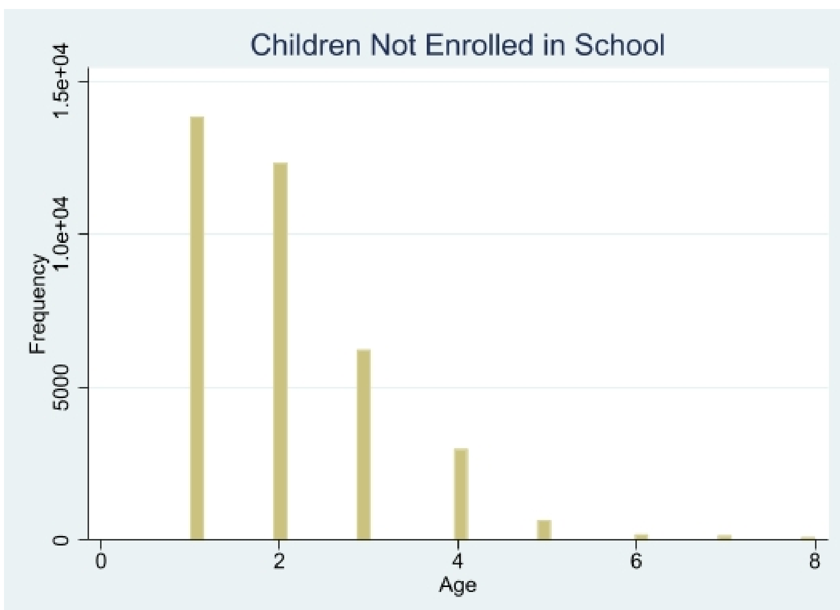

Table 1: WIC participation

Child's Age	WIC Recipient	Total	Proportion
1	365	1,999	0.183
2	288	1744	0.165
3	159	982	0.162
4	74	444	0.167
5	22	167	0.132
6	1	65	0.015
7	6	52	0.115
8	4	29	0.138

The incredibly low sample size after the age of 4 is very concerning and indicates the a fuzzy RD design is not the best choice. We also see sudden drops in our sample size at ages 3, 4, and 5 as children begin to enroll in school and are excluded from our sample. This raises concerns that the mothers that don't enroll their older children in school are different from the mothers of younger children that are not enrolled in school. To ensure balance across the groups above and below the cutoff, we compare distributions of key characteristics, with results as follows:

Table 2: Summary Statistics

	$Age_{child} < 5$	$Age_{child} \geq 5$	Difference
Observations	9,337	313	
WIC participation	0.171	0.105	-0.066***
	[0.377]	[0.307]	(0.022)
Attending School	0.105	0.046	-0.060**
	[0.307]	[0.210]	(0.003)
Number of Children	2.151	2.003	-0.147**
	[1.142]	[1.108]	(0.066)
Family Size Excluding Children	1.988	1.923	-0.064
	[1.703]	[0.665]	(0.040)
Black	0.106	0.128	0.002
	[0.307]	[0.334]	(0.018)
Hispanic	0.200	0.201	0.001
	[0.400]	[0.402]	(0.023)
Living with a Parent or Partner	0.462	0.469	0.007
	[0.498]	[0.500]	(0.029)
Any Disability	0.029	0.032	0.003
	[0.167]	[0.175]	(0.012)

We find significant differences in WIC participation between mother of children below the eligibility threshold and mothers of children above the eligibility threshold. This confirms our decision to use a fuzzy RDD.

We also find significant differences in the share of mothers attending school as well as their number of children. Therefore, we include controls for these variables as indicated in equations 4 and 5 by G_i. Due to endogeneity concerns, but we also include fixed effects for state of residence, year, measure of food insecurity (excluded when outcome variable of interest is food insecurity), whether a mother is black and family income. These factors may impact WIC eligibility and participation as well as our outcome variables of interest and for this reason, we include them in our controls.

5 Results

The small sample size of mothers with children after the age of 4 is incredibly concerning and means that our results should not be accepted. Further concerns diminishing the validity of the results are included in the discussion section. Instead, these results can be thought of as suggestions for areas of future research. Under the assumption that the small sample size is sufficient, we can measure the effect of exit from WIC on the maternal labour market and food insecurity.

Using our first stage regression, we measure the effect of WIC participation as shown below:

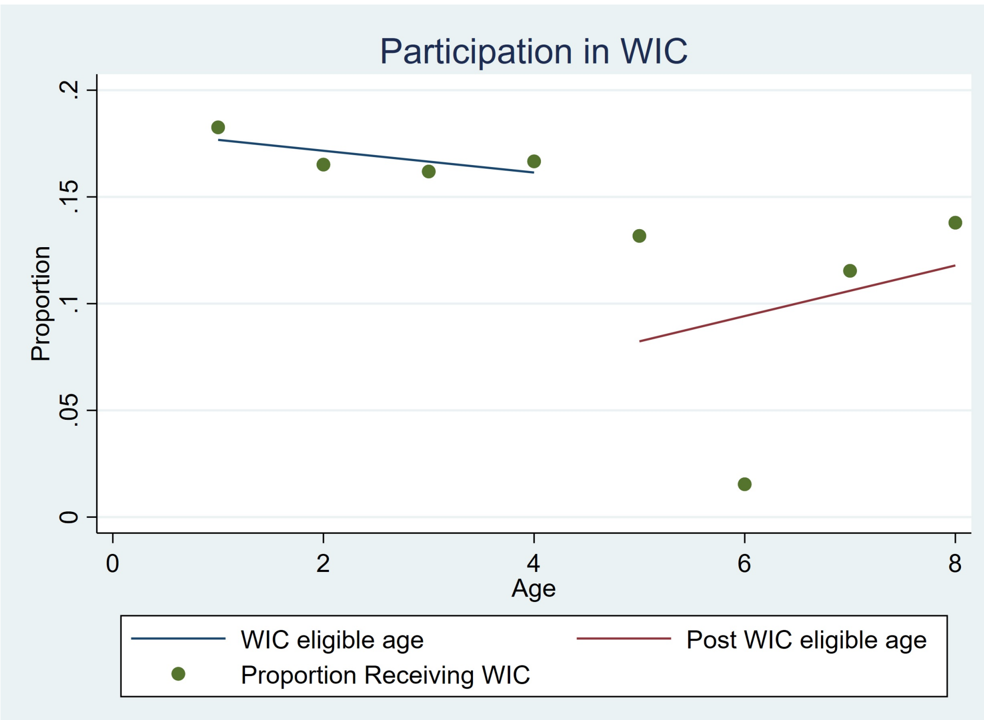

Figure 1: Proportion of WIC participation given age of child

The non-linearity of WIC participation is likely due to the rapidly shrinking sample size past the age threshold of 5. We use this this trait when constructing our first stage regression equations.

The results of our first stage equation (2), are constructed with a probit model and shown below:

Table 3: 1st Stage - WIC with Food Insecurity Controls

VARIABLES	(1) Model 1	(2) Model 2
Age Eligible	0.362*	0.739**
	(0.220)	(0.318)
Age	0.0131	0.113*
	(0.0369)	(0.0658)
Age^2		0.0329*
		(0.0179)
Number of own children in household	-0.0151	-0.0151
	(0.0335)	(0.0335)
Attending School	0.0558	0.0519
	(0.112)	(0.112)
$Black_{mom}$	-0.181	-0.171
	(0.122)	(0.123)
Observations	2,082	2,082
Flexible Slopes	YES	YES
Non-Linear Trends	NO	YES

Robust standard errors in parentheses
*** p<0.01 ** p<0.05 * p<0.1

The results of our first stage regressions indicate that our constructed dummy, $Age_{eligible}$ significantly predicts participation in WIC among households living below 185% poverty levels. Adding non-linear trends, as in Model 2, more accurately captures WIC participation.

We use the predicted WIC values generated from the above regression to measure the effect of exit from WIC on maternal employment and hourly wage.

Table 4: 2nd Stage - Employment

VARIABLES	(1) Model 1	(2) Model 1	(3) Model 2
$\hat{WIC}$	0.777	0.624	1.047*
	(0.551)	(0.721)	(0.634)
Age	0.0588	0.0613	0.0252
	(0.0400)	(0.0400)	(0.0479)
Age^2			-0.0145
			(0.0153)
Age $\times$ $\hat{WIC}$		-0.0714	-0.846*
		(0.290)	(0.440)
Age^2 $\times$ $\hat{WIC}$			-0.352**
			(0.166)
Number of own children in household	-0.0299	-0.0304	-0.0307
	(0.0403)	(0.0402)	(0.0406)
Attending School	0.405***	0.406***	0.398***
	(0.146)	(0.146)	(0.147)
Black$_{mom}$	-0.0533	-0.0552	-0.0483
	(0.117)	(0.116)	(0.117)
Observations	2 395	2 395	2 395
Flexible Slopes	NO	YES	YES
Non-Linear Trends	NO	NO	YES

Robust standard errors in parentheses
*** p<0.01 ** p<0.05 * p<0.1

Using our second stage regression, we predict whether a mother is employed using a probit model. The results are shown above and indicate that exit from WIC fails to have a significant effect on whether a mother is employed. However, other controls do significantly affect a mother's employment status. The sign of these coefficient suggests that mothers attending school are more likely to be employed.

Table 5: 2nd Stage - Hourly Wage

VARIABLES	(1) Model 1	(2) Model 2	(3) Model 3
$\hat{WIC}$	-4.195	-4.198	-4.454
	(2.754)	(3.035)	(3.298)
Age	-0.0150	-0.0157	-0.161
	(0.245)	(0.253)	(0.253)
Age^2			-0.0692
			(0.109)
Age $\times$ $\hat{WIC}$		-0.174	-0.219
		(1.111)	(1.917)
Age^2 $\times$ $\hat{WIC}$			-0.0113
			(0.946)
Number of own children in household	-0.176	-0.155	-0.166
	(0.389)	(0.392)	(0.397)
Attending School	-0.279	-0.266	-0.323
	(1.011)	(1.020)	(1.032)
Black$_{mom}$		-0.573	-0.550
		(0.735)	(0.749)
Observations	426	426	426
R-squared	0.476	0.477	0.477
Flexible Slopes	NO	YES	YES
Non-Linear Trends	NO	NO	YES

Robust standard errors in parentheses
*** p<0.01 ** p<0.05 * p<0.1

Using the same predicted WIC values, but now controlling for maternal employment status, we measure the effect of WIC on our second outcome variable of interest, hourly wage. Our results, shown above, indicate that WIC participation fails to significantly affect changes in hourly wage.

For our third outcome variable of interest, whether a household is very food insecure, we exclude food insecurity as a control. This results in a different first stage regressions resulting from a probit model. Our results are shown below:

Table 6: 1st Stage - WIC without Food Insecurity Controls

	(1)	(2)
VARIABLES	Model 1	Model 2
Age Eligible	0.329	0.703**
	(0.219)	(0.321)
Age	0.00903	0.109*
	(0.0368)	(0.0656)
Age^2		0.0329*
		(0.0179)
Number of own children in household	-0.0208	-0.0209
	(0.0334)	(0.0334)
Attending School	0.0521	0.0490
	(0.112)	(0.112)
Observations	2,082	2,082

Robust standard errors in parentheses
*** p<0.01 ** p<0.05 * p<0.1

The results of our first stage regression on food insecurity indicate that without controls for food insecurity, age eligibility fails to accurately predict WIC participation unless non-linear trends are included.

Table 7: 2nd Stage - Household Very Food Insecure

	(1)	(2)	(3)
VARIABLES	Model 1	Model 2	Model 3
$\hat{WIC}$	0.0208	0.736	0.538
	(0.289)	(0.703)	(0.480)
Age	0.0185	0.0171	0.00643
	(0.0236)	(0.0239)	(0.0332)
Age^2			0.00446
			(0.0112)
Age $\times$ $\hat{WIC}$		-0.0180	
		(0.0480)	
Age^2 $\times$ $\hat{WIC}$			-0.0773
			(0.0518)
Number of own children in household	0.00133	0.00137	-0.0264
	(0.0260)	(0.0260)	(0.0294)
Attending School	-0.109	-0.109	-0.171
	(0.0956)	(0.0956)	(0.108)
Black$_{mom}$	0.0983	0.0982	0.147
	(0.0854)	(0.0854)	(0.0997)
Observations	3,952	3,952	3,941
Flexible Slopes	NO	YES	YES
Non-Linear Trends	NO	NO	YES

Robust standard errors in parentheses
*** p<0.01 ** p<0.05 * p<0.1

Our third outcome variable of interest, whether a household is very food insecure, has an insignificant relationship with WIC participation as well as our other controls. This may reflect WIC's nature as a food supplement. Food packages are targeted to individuals in the family, and not aimed at the household as a whole. The CPS marks a household as very food insecure if any individual in the household changed their eating patterns because of a lack of resources. This variable does not encompass children, specifically. For this reason, it is important to look at food security among children.

Table 8: 2nd Stage - Children Very Food Insecure

VARIABLES	(1) Model 1	(2) Model 2	(3) Model 3
$\hat{WIC}$	0.0577	-6.244***	-7.029***
	(0.681)	(0.924)	(1.382)
Age	0.145**	0.155***	0.0790
	(0.0582)	(0.0584)	(0.0671)
Age^2			-0.0191
			(0.0278)
Age$\times$ $\hat{WIC}$		0.762***	
		(0.129)	
Age^2 $\times$ $\hat{WIC}$			0.782***
			(0.180)
Number of own children in household	-0.0312	-0.0296	0.0246
	(0.0740)	(0.0742)	(0.0820)
Attending School	0.355	0.366	0.179
	(0.248)	(0.248)	(0.282)
Black$_{mom}$	0.0754	0.0883	-0.249
	(0.196)	(0.200)	(0.241)
Observations	817	817	814
Flexible Slopes	NO	YES	YES
Non-Linear Trends	NO	NO	YES

Robust standard errors in parentheses
*** p<0.01 ** p<0.05 * p<0.1

Our final outcome variable of interest, whether children in the house are very food insecure does have a significant relationship with food insecurity. Our interaction term between age and predicted WIC values from our third model was dropped for co-linearity, but the results show that WIC participation significantly alleviates whether children in a household are very food insecure.

6 Discussions and Conclusions

Our results suggest that exit from WIC does not impact maternal employment status, wage, or whether a household is very food insecure. Exit from WIC does alleviate food insecurity in children. However, our methodology faces flaws which are discussed below.

One flaw in our design is the sharp decline in sample size past the age eligibility threshold. The mothers of older children may act as outliers, overly influencing the trend of the cohort above the age eligibility threshold. In addition, states vary in their school attendance requirements and mothers with children unenrolled in school may experience extenuating circumstances that are not captured by CPS data. For example, children may be exempt from school attendance requirements in case of severe medical conditions or disabilities. The CPS does not collect data on whether children under the age of 15 has a disability and therefore our study could not control for the presence of childhood medical conditions. Therefore, children with medical conditions may be over represented in the cohort above the age eligibility threshold. These extenuating circumstances may also influence food security and would lead to underestimating the effect of WIC on food security.

The CPS data also fails to indicate the type of WIC package received by the household. WIC food packages can be for pregnant women, postpartum women, infants, or children. We would expect most of the food packages, especially below the age eligibility threshold, to reflect child food packages. However, some of the food packages, especially above the age eligibility cutoff, may be for pregnant women. The CPS does not include data on whether a mother is pregnant, and these food packages would confound and potentially diminish the impact of the WIC child food packages.

Some policy makers are concerned that providing food benefits will lead to changes in labour force participation along the external margin. This paper supports existing literature that providing food benefits will not lead to such changes. Especially during a pandemic, providing food benefits supports a very vulnerable population when many firms are not hiring.

7 References

- Cebula, R. J., Coombs, C. K. (2008). Recent evidence on factors influencing the female labor force participation rate. Journal of Labor Research, 29(3), 272-284.
- Herman, D. R., Harrison, G. G., Afifi, A. A., Jenks, E. (2004). The effect of the WIC program on food security status of pregnant, first-time participants. Family Economics and Nutrition Review, 16(1), 21.
- Huffman, S. K., Jensen, H. H. (2008). Food assistance programs and outcomes in the context of welfare reform. Social Science Quarterly, 89(1), 95-115.
- Kreider, B., Pepper, J. V., Roy, M. (2016). Identifying the effects of WIC on food insecurity among infants and children. Southern Economic Journal, 82(4), 1106-1122.
- https://drexel.edu/hunger-free-center/research/briefs-and-reports/WIC-improves-child-health-outcomes/
- https://www.fns.usda.gov/wic/participant-and-program-characteristics-2018-charts
- WIC Participant and Program Characteristics 2018: Final Report
- National WIC Association,WIC: Solid Returns on Investment While Reducing the Deficit
- Seth, J. G., Isbell, M. G., Atwood, R. D., Ray, T. C. (2015). Addressing language barriers in client-centered health promotion: Lessons learned and promising practices from Texas WIC. Health promotion practice, 16(3), 320-328.
- Moffitt, R. (1983). An economic model of welfare stigma. The American Economic Review, 73(5), 1023-1035.
- Besley, T., Coate, S. (1992). Understanding welfare stigma: Taxpayer resentment and statistical discrimination. Journal of Public Economics, 48(2), 165-183.
- Pak, T. Y. (2020). Welfare stigma as a risk factor for major depressive disorder: Evidence from the Supplemental Nutrition Assistance Program. Journal of affective disorders, 260, 53-60.
- Liu, C. H., Liu, H. (2016). Concerns and Structural Barriers Associated with WIC Participation among WICâEligible Women. Public Health Nursing, 33(5), 395-402.
- Brewer, M., Duncan, A., Shephard, A., Suarez, M. J. (2006). Did working families' tax credit work? The impact of in-work support on labour supply in Great Britain. Labour economics, 13(6), 699-720.

- East, C. N. (2018). Immigrantsâ labor supply response to Food Stamp access. Labour Economics, 51, 202-226.
- Meyer, B. D., Rosenbaum, D. T. (2001). Welfare, the earned income tax credit, and the labor supply of single mothers. The quarterly journal of economics, 116(3), 1063-1114.
- Coile, C., Gruber, J. (2001). Social security incentives for retirement. In Themes in the Economics of Aging (pp. 311-354). University of Chicago Press.

Who Chooses to Work for Themselves? Cultural Integration and Incentives for Immigrant Self-Employment in the UK *

Merton Ngan Timothy Wong Amy Li Bryan Ling Wei Guang Tan Max Schachermayer

August, 2021

Abstract

We examine the effect of cultural integration to the UK on the probability of self-employment and potential wages via citizenship status, by exploiting self-selection in employment and self-employment choice to compare potential wages in employment choice to self-employed or employed sectors. Using two waves of panel data from the UK Household Longitudinal Survey in 2009 and 2014, we use variation in time in conjunction with the Heckman two-stage selection model to correct for the problem where people endogenously self-select into employment/self-employment, which leads to biased and inconsistent estimates. We then explore the breakdown of these effects by self-employment and employment sectors. We find that there is a negatively significant effect of national identity on the probability of self-employment and a significant effect of being an immediate immigrant (second generation immigrant) for employed workers only. We make further extensions to account for individual unobservable heterogeneity and attempt to check for robustness in other factors such as parental education. Our results corroborate with findings in the literature, indeed cultural identity plays a significant role in determining the sector of employment and potential wages.

---

*Department of Economics, London School of Economics and Political Science, London, UK. Corresponding email address: p.ngan@lse.ac.uk. We are grateful to Judith Shapiro for her invaluable help in all stages and motivational support. We thank Mikael Johansson, Stefanus Phan, Eddy Zou, Mina Burns and Jason Jia for helpful discussions and comments.

1 Introduction

The significant inflows of immigrants into Western countries in the post-war era has generated considerable research interest into the economic status and contributions of immigrants. Immigrants are widely perceived to be highly entrepreneurial: self-employment and business ownership not only improve labour market assimilation and economic outcome of immigrants, but also inject a dose of dynamism to the host economy via improving human capital.

Traditionally, immigrants in the UK, particularly those who are an ethnic minority, are perceived to be associated with low skill and low return sectors. However, recently, increased net migration, diversity in the country of origin, along with changing sources of migration (asylum seekers and economic migration), means that over generations, the composition of skill level and socio-economic status in migrants has changed over time. This gives a compelling reason to study not only the effect of being an immigrant on the self-employment choice and potential wages, but also how this effect changes over time and across generations. In the UK, immigrants, especially ethnic minorities are more likely to be unemployed, while those employed tend to earn less than their native counterparts (Nandi and Platt (2015)). The reason for this varies in the literature, with the main reasons being cultural integration, language proficiency, and proximity to other immigrants. This paper aims to assess the role of cultural integration as a primary driver of the self-employment choice, where the previous literature suggests citizenship take-up as a good measure of the strength of an immigrant's connection to the country of arrival and cultural integration.

To account for endogeneity in the self-employment choice, we estimate models for the self-employment choice and use the resulting estimated Inverse Mills ratio to determine the potential wages of the employed and self-employed.

The use of cultural integration as the primary driver of the self-employment decision is not one commonly found in the literature pertaining to the UK. This paper contributes to the literature by using recent micro-data to provide the latest evidence on the 'immigrant effect' for the self-employment decision and wage differentials in the UK. In particular, we apply the Heckman style regression models, while exploiting the UK Household Longitudinal Survey's wide range of questions relating to migration and cultural integration. In using the data provided in the survey, we will also critically analyse the strengths and limitations of the dataset.

We find that there is a negative and statistically significant impact of cultural integration on the probability of being self-employed and wages. Our sensitivity analysis suggests that the effects vary in significance and magnitude, largely reflecting sample size limitations rather than confounding.

The rest of this paper is organised as follows: Section 2 reviews the existing literature, Section 3 describes the data sources used, Section 4 outlines the main empirical strategy, Section 5 reviews the results, Section 6 reviews sensitivity and robustness tests conducted and Section 7 concludes.

2 Literature Review

There exists a large body of international literature on the determinants of immigrant self-employment; the topic has garnered interest from both economists and sociologists, with research ranging from government-commissioned surveys to bespoke case studies. Currently, there is relatively more abundant evidence on US ethnic entrepreneurship, and scarcer on UK and the rest of Europe. In this section, we focus on previous literature that seek to answer why natives and immigrants have differential propensities to self-employment as well as phenomena in the labour market.

A prominent approach to understanding the higher propensity amongst immigrants to become self-employed is the ethnic enclave theory. The tendency of immigrant communities to concentrate in select geographical areas creates de facto "protected markets", as Cochran (1972) proposes: certain cultural preferences and specialised needs of ethnic minorities are best catered for by co-ethnic businesses. Various sociological case studies on specific ethnic groups, such as Wilson and Portes (1980) on Cuban enclaves in Miami, or Aldrich et al. (1985) on Asian populations in Bradford and Leicester, provide substantial support to the hypothesis that ethnic residential concentration explains patterns of business patronage.

Closely allied to the ethnic enclave theory is the language proficiency theory. It is argued that immigrants who are less fluent in the host country language face limited labour market opportunities and are thus more likely to be driven into self-employment in ethnic businesses, where there are no linguistic barriers. Evans and Leighton (1989) extends upon the "protected market" hypothesis discussed above by proposing that a linguistically isolated labour pool of workers can be more profitably tapped by co-ethnic entrepreneurs, who can readily communicate with them. The evidence from Australia does indeed confirm that the higher the proportion of adults who are not fluent in English within an ethnic community, the higher the likelihood that they will go on to set up their own businesses. This is largely consistent with the later Australian literature by Breunig et al. (2013), who find that wage assimilation occurs most slowly for immigrants from non-English speaking backgrounds. However, both US-and UK-based studies, including Portes and Zhou (1993) as well as Clark and Drinkwater (2000), seem to conclude the opposite - that those with poorer English skills have lower probabilities of becoming self-employed.

Language proficiency is an integral aspect of human capital, which is expected to improve the longer the immigrant has resided in the host country. The related immigrant status theory considers the number of years elapsed since the immigrant's first arrival, closely associated with the accumulation of several forms of capital: including human, physical and capital. Friedberg (2000) highlights the close relationship between home-country human capital accumulation and the types of jobs available to the immigrants in the host-country, though the transferability of home-country human capital and experience is low, hindering employment prospects for immigrants upon first arrival in the host-country.

With increasing immigration and changing ethnic composition of migrants, society is now increasingly multicultural with mixed consequences. In the UK, multiculturalism is perceived to have connections to separatism and a lack of desire to conform to core national values. This issue has been particularly prevalent for second generation immigrants. However, Bartram (2021) indicates that in the UK and Europe, immigrant ethnic minorities tend to develop a British national identity similar to that of white natives. This implies that ethnic identity is not necessarily separate from national identity and thus multiculturalism is not a threat to national identity. However, the strength of one's national identity and propensity to be culturally integrated depends positively on factors such as socio-economic opportunity, prevalence of discrimination and the aforementioned factors. Since national identity is not directly measurable, it is pivotal to determine a valid proxy for it to use for regression analysis. Previous literature indicates that those who take-up citizenship increase attachment to British identity relative to that of non-citizens, and thus can be used as a proxy for national identity. However, practically, there are many other factors that affect citizenship take-up outside national identity, including political climate, income and perceived social and economic opportunities. Hamilton (2000) uses three categories to explain the earnings gap between self-employed and employed wages: investment and agency models which argue that differences in earning profiles across sectors account for the earnings gap; matching and learning models which argue that heterogeneity in sector-specific unobservables results in self-selection into self-employment or paid employment; variation in working conditions across sectors where entrepreneurs may sacrifice earnings for benefits in self-employment. He finds that self-employed workers tend to have lower initial earnings and earnings growth relative to comparable employed workers, and median self-employed wages are always lower than on entry level employed wages.

Borjas (1986) finds that in the US, immigrants are more likely to be self-employed due to the ethnic enclave theory. Immigrants have higher self-employment rates than natives in North America. Self-employed immigrants are significantly more likely to be in a "retail job" industry compared to self-employed natives, indicating immigrants assimilate by opening small shops. In contrast to Hamilton's findings, self-employed immigrants have higher annual salaries than employed workers and have higher salaries than other self-employed natives but attributes it to returns on physical capital of self-employed workers.

Clark et al. (1998) finds that employed wages are lower for immigrants compared to self-employed wages, and the wage disparity increases over time. Ethnic minorities also have a higher propensity to be self employed compared to whites but there is significant variation across ethnic groups. In addition, ethnic minorities earn less in employed work than white workers, and an increase in the earnings gap for ethnic minorities over time is in line with a shift into self-employment. This suggests that self-employment is a way out for immigrants to escape discrimination in the employed sectors in the UK.

3 Data

Our primary data sources include waves 1 and 6 in the UK Household Longitudinal Survey (UKHLS). The UKHLS is a panel survey that asks over 200 questions on topics such as employment, income, family, health and civic participation. In addition to the main sample of around 10,000 individuals, there was a booster sample of over 6,000 adults in wave 6. A key distinguishing feature of this survey is the breadth of topics asked, such as the number of times someone had been assaulted or harassed due to their ethnicity, job satisfaction, detailed educational qualifications, membership in workplace unions, political affiliations and so on. This allows us to, as a potential extension for future research, form a well-rounded understanding of the sample and thus find more nuanced counterfactuals. Further, using a panel dataset allows us to track the same individuals over time, accounting for individual-level time-invariant confounders (through the inclusion of individual fixed effects), and use variation in reported national identity over time to estimate the desired effect.

Nevertheless, there are two major limitations of this data source: first, our empirical strategy relies on variation in responses over time to work, but most variables we used had a significant lack of variation over time. This forced us to remove a significant number of observations. We recognise that there this leads to a trade-off between exacerbated problems of self-selection on unobservables. We also reduced the number of waves of survey used, which helps alleviate potential time series complications when using longer panels. However, this reduced the number of potential observations.

Response rate varied heavily by question and immigrant status. For some questions response rate was too low, so their inclusion in the study as a potential control was ruled out. For certain variables which were pivotal to our study, we had to impute observations that didn't respond to at least one of these questions as well as use common sense for time-invariant variables.

This left the study with 7573 working-age observations (between the ages of 16-60). 63.62% are born in the UK. 13.34% are self-employed, which is lower than preferred, but is large enough to carry on analysis. Of the immigrants, 13.94% are self employed, which is only marginally higher than the proportion of natives who are self employed at 12.99%, which seems contrary to the belief that immigrants are disproportionately self-employed. The gender split is 50.6% female and 49.4% male. Figure 1a summarizes the ethnic group distribution, which is predominantly white as expected and is similar for both self-employed and employed sectors. Detailed breakdowns of distributions of key categorical variables and wages broken down by employment sector are shown in Figures 1 to 4 attached in the paper.

In line with the ethnic enclave theory, we included controls for government office regions. It is interesting to note that a large majority of immigrants tend to live in London, more than natives in this sample, which reflects larger earning opportunities in the city. With a special license, we could narrow down these regions significantly but due to time constraints we regrettably were not able to do so. Therefore, we are on the side of caution when interpreting results for this particular variable as government office regions are large and are not necessarily representative.

Wage distributions appear to be less right-skewed, there is a lower variance for the self-employed sector than the employed sector, but the latter may be due to low sample size. We also find a number of negative wage values which inspires us to transform the wage distribution by inverse hyperbolic sine.

We are also interested in whether an individual being a descendent of an immigrant family with deeper roots in the UK affects their rate of cultural assimilation and thus their potential wage outcomes. Using the available data, we can examine whether an individual is a first, second, third, fourth, or longer generation immigrant. The distribution of respondents in terms of being n-th generation is shown in Figure 1. Overall, there is some heterogeneity that we can use to investigate this in the main regression.

From Table 3(a), we find that most of the sample's highest educational attainment is an undergraduate degree regardless of employment sector. A larger proportion of immigrants have at least a degree, and across all generations of immigrants, a majority of them have at least a degree, but the share increases the more recent the arrival. This supports prior literature showing increasing education and socio-economic status of immigrants to the UK.

Parents' education is a good indicator of socio-economic status and potentially may affect the line of work of their children, the distribution of which is shown in Figure 4. We find that fathers tend to be less educated, more of whom only have a high-school level education. Unfortunately, updated and well-responded to data of parents' occupation is not readily available, which is also a good indicator of socio-economic status.

Given that we have panel data at our disposal, we are able to track changes in the same observation over time. Table 3(b) shows the observation counts of those who switch from the employed sector to the self-employed sector and vice versa. A majority of natives that switch from employment to self-employment are 2nd generation immigrants. Whlie we would like to know if an increase in instances of xenophobic related abuse will affect propensity of self employment, but response rates for those questions are low and lack variation over time.

4 Identification Strategy

A key threat to identifying the main causal effects of interest in our analysis is the self-selection problem in the self-employment decision. As discussed in the previous Literature

Review section, we choose to focus on the matching and learning models as the base model to explain differences in propensity to be self-employed and subsequent wage differences. In line with predecessors' strategies, we use a Heckman two-stage correction model to correct for this self-selection, which in theory causes biased and inconsistent estimates.

Theoretically, the model filters out endogeneity due to the sample selection problem under the condition that error terms in the wage models are jointly normal (and thus correlated) with zero mean and unconstrained covariance structure (Heckman (1979)), explanatory variables in the selection and main/outcome equations are uncorrelated with their respective error terms as well as usual rank conditions to develop Generalised Least Squares (GLS) estimators.

In line with Roy's searching/compensation model, we believe workers will self-select into self-employment if wages are larger in that particular sector. While this is not as nuanced a view as other hypotheses that account for entrepreneurs sacrificing higher potential earnings for autonomy in the workplace, this simplification does not detract from the intention of the analysis. After correcting for the self-selection, we then use the following wage estimation equations for the self-employed and employed as a simple linear model with relevant controls and the Inverse Mills Ratio derived from the first stage. We define the choice to be self-employed, represented in the model by L_i, using a binary variable that represents whether the respondent is self-employed or not. The value of this binary variable is defined by the latent variable L_{it}^*:

$$L_{it} = \begin{cases} 1, & L_{it}^* \geq 0 \\ 0, & Otherwise \end{cases} \tag{1}$$

Our model assumes that $L_{it}^* = \alpha_0 + \alpha_1 D_{it} + \alpha_2 X_{it} + \alpha_3 \gamma_t + \epsilon_i$. The first stage then uses this model of self-employment to generate a probit model to estimate the probability of being self-employed for individual i in wave t ($L_i = 1$):

$$\mathbb{P}(L_{it} = 1 | X_{it}, D_{it}) = \Phi(\alpha_0 + \alpha_1 D_{it} + \alpha_2 X_{it} + \alpha_3 \gamma_t) \tag{2}$$

Where Φ is the standard normal CDF, D_{it} is each respondent's strength of British identity, γ_t is the wave fixed effect, and X_{it} is the set of other covariates, including, age, education background, etc. It is clear here that the error term ϵ_i must be (assumed to be) normally distributed to apply the normal cumulative distribution.

From the first stage results, we can also derive the Inverse Mills Ratio λ_{it} as given below using the estimated value for αs by Maximum Likelihood Estimation (MLE) on the first stage:

$$\lambda_{it} = \frac{\Phi(\alpha_0 + \alpha_1 D_{it} + \alpha_2 X_{it} + \alpha_3 \gamma_t)}{\phi(\alpha_0 + \alpha_1 D_{it} + \alpha_2 X_{it} + \alpha_3 \gamma_t)} \tag{3}$$

where Φ is the standard normal cdf, ϕ is the standard normal pdf.

We transform wages to account for the fact that self-employed wages can be negative due to losses suffered in early stages of operating a firm and the large exponential-like gap in wages by employing Johnson (1949)'s inverse hyperbolic sine transformation on all wages (denoted by $IHS(Y_{it})$). This also allows us to work with a more condensed wage distribution and models inequality better. We next consider the following model in the second stage:

$$\begin{cases} IHS(Y_{it}^{(0)}) = \beta_0^{(0)} + \beta_1^{(0)} G_{it} + \beta_2^{(0)} Z_{it} + \beta_3^{(0)} \gamma_t + v_{it}^{(0)}, & L_{it} = 0 \\ IHS(Y_{it}^{(1)}) = \beta_0^{(1)} + \beta_1^{(1)} G_{it} + \beta_2^{(1)} Z_{it} + \beta_3^{(1)} \gamma_t + v_{it}^{(1)}, & L_{it} = 1 \end{cases} \tag{4}$$

where Y_{it} is the wage, G_{it} is the set of immigrant generation dummies, Z_{it} is the set of relevant covariates, and γ_t is the wave fixed effects. It is worth noting that the second stage for the wage of the employed can only be estimated using the subset of the data who are employed, and vice versa for the second stage of the self-employed.

Applying the Heckman selection model outlined above, we can incorporate the Inverse Mills Ratio from the first stage via the conditional expectations of the second stage:

$$\begin{cases} \mathbb{E}(IHS(Y_{it}^{(0)})|G_{it}, Z_{it}, \gamma_t, L_i) = \beta_0^{(0)} + \beta_1^{(0)} G_{it} + \beta_2^{(0)} Z_{it} + \beta_3^{(0)} \gamma_t + \rho^{(0)} \sigma^{(0)} \lambda_i, & L_i = 0 \\ \mathbb{E}(IHS(Y_{it}^{(1)})|G_{it}, Z_{it}, \gamma_t, L_i) = \beta_0^{(1)} + \beta_1^{(1)} G_{it} + \beta_2^{(1)} Z_{it} + \beta_3^{(1)} \gamma_t + \rho^{(1)} \sigma^{(1)} \lambda_i, & L_i = 1 \end{cases} \tag{5}$$

where ρ is the correlation coefficient between the v_{it} and ϵi, σ is the standard deviation of v_i. The Inverse Mills Ratio seeks to address the selection issue that is central to our study. We use the estimated IMR in the second stage, which allows us to determine the effect of different key variables on wages and the propensity to be self-employed via MLE.

5 Results

As shown in Table 4, we find that the probability of being self-employed decreases with British identity and this effect is statistically significant. This may be due to a greater sense of community that comes with stronger national identity and thus workers may not feel the need to open their own businesses in line with earlier literature such as ethnic enclave or language proficiency theory.

We find a number of statistically significant effects in other control variables: a negative effect of future financial situation which may be due to the volatility and lower earnings potential of self-employed work and an unwillingness to undertake higher risk with an expectation of higher future earnings; job satisfaction has a positive effect; number of children has a positive effect, which seems counter-intuitive due to higher volatility and lower earnings potential in being self-employed; age also has a positive effect which may be due to a desire for people to pursue their interests in their later years. Interestingly, the effect of the current financial situation is not statistically significant. The IMR is significant only for the employed sector and is negative. As with other results for the self-employed sector, we interpret the non-significant effect of the IMR with caution, given the low sample size. The sign on the IMR tells us the direction of selection: surprisingly, there is negative selection on employed wages, but positive selection on self-employed wages. We would have expected positive selection on both as people choose to self-select into a particular sector if they are more suited to it.

Interpreting results from a nonlinear model is not as straightforward as a marginal effect interpretation from a linear model. We also exercise caution in interpreting size of effects as they appear to be drastically different for employed and self-employed due to large volatility in estimates caused by low sample size.

As stated previously, we model wages separately for employed and self-employed workers. We find that the effect of being an immigrant is significant only for a second generation immigrant for employed workers. We find that the effect is marginally positive for second generation and third generation immigrants for employed workers, second generation and fourth+ generation for self employed workers. Therefore, under our identification assumptions, being an immigrant increases wages for both self-employed and employed workers, which is similar to the results of Borjas (1986). We find that effects in the second stage are mostly statistically significant for the employed but are not for the self-employed. This is likely due to a low observation count for self-employed people that makes estimates imprecise.

We find that being a part-time worker has a negative effect on wages, which is expected due to lower earnings potential of part-time work. This effect is larger for self-employed workers, perhaps again due to lower earnings potential of self-employed workers, which is even lower in the case of part-time workers in loss-making firms.

We then break down effects on earnings based on ethnic group, using the White English/Scottish/Welsh/Northern Irish Group as our reference group. We find that the ethnicity effect on wages for the employed is significantly smaller for a large number of ethnicities as shown in Figure 5. Employers of African (and other black), Indian, Bangladeshi, Pakistani and other Asian descent have a significantly lower effect (using a 95% confidence interval) than Whites, all else equal. This corroborates with a well-documented ethnic discrimination factor in employed sectors, which is larger for Pakistanis and those of African descent than other groups. We believe this effect may be different across sectors and remains an interesting area for future research. There is no group that has significantly higher earnings than Whites in the employed sector, but the standard errors are fairly large for other ethnic groups, making estimates of the earnings gap difficult to ascertain.

For self-employed workers, the standard errors are larger due to lower observation count across the board, so conclusions drawn from these results are difficult to interpret. We cannot ascertain an ethnic discrimination factor for the self-employed sector. However, we find that only workers of Other Mixed and White & Black Caribbean descent have significantly smaller earnings than Whites, but no group has significantly larger earnings than Whites.

We also break down effects by educational attainment, which are relative to that for those whose highest educational qualification is a (university) degree. We find that for employed workers, having any other highest educational qualification level leads to significantly lower earnings relative to those who only have a degree. Those with no qualifications suffer the most, followed by other qualifications, which is expected as education improves human capital by imparting knowledge and soft skills and improving earnings potential by virtue of signalling. Those with high school earnings suffer less and those with only GCSEs suffer more than those with A-levels, which corroborates with the signalling theory.

Interestingly, those with other higher qualifications have a smaller earnings effect. This may be due to over-qualification, which can reduce earnings potential. For the self-employed, there are no educational attainment levels that are significantly better than having a degree. Again, this is due to large standard errors caused by low observation count. We cannot say with confidence that there are returns to education in the self-employed sector. There are clearly positive returns to education for employed workers.

Overall, we find that there are a number of factors contributing to the self-employment decision, including British identity. However, we find that being an immigrant or not doesn't significantly improve earnings outcomes unless they are a second generation immigrant in the employed sector and several other well-documented phenomena of earnings in the workplace are corroborated. The directions of effects are fairly similar for both employed and self-employed, which indicates that most factors determining wages are common to both self-employed and employed. Our analysis was limited by data constraints for self-employed workers. The large standard errors and lack of variation in their education and ethnicities, lead to imprecise estimates.

6 Robustness Checks

We conduct several variations to our previous model to check for robustness in our results to specification choice.

We attempt to examine if the effect of being an immigrant is robust to using a broader definition of immigrant by replacing the generation variable with the *ukborn* variable, a dummy variable indicating if the individual was born in the UK, the results are shown in Table 5.

The generation variable breaks down each person by the generation of immigrant they are, which is either 1st, 2nd, 3rd, or 4th+. The *ukborn* variable classifies 2nd, 3rd and 4th+ generation immigrants as natives and 1st generation immigrants as non-natives. We find that being an immigrant has a significant effect on wages in the employed sector but not in the self-employed sector, which is a relatively similar result as that from the main regression. Effects of other controls are also similar to that of the main regression, indicating that our estimates and their interpretations are fairly robust to the narrower definition of immigrant.

To account for the effect of parents' education on the probability of self-employment, we added these control variables to the first stage of our model, results are shown in Table 6. We find that the effect of the current financial situation is now significant, which may perhaps be due to higher financial security incentivising people to "work for themselves" and have more autonomy over their work. The second stage results are largely the same, but due to low sample size in the self-employed sector results appear to be somewhat unreliable. The most notable difference is that the IMR in both employed and self-employed sectors are now both not significant, which implies that there is no evidence of self-selection, but this is likely due to low sample size.

We then run specifications adding a dummy variable indicating whether or not an individual's first job is self employed to the first stage regression, results of which are shown in Table 7.

We expect this effect to be significant and other effects to be robust to this addition. We find that in the first stage, the effect of current financial situation on the probability of self-employment is now significant and positive for the reasons stated above. Effects for all generations in both the employed and self-employed sectors are now not significant in the second stage, which seems counter-intuitive as we would expect being a native would improve wages in accordance with other literature as previously discussed. Other effects remain qualitatively unchanged. However, the IMR for both groups is now not significant which is likely due to small sample size as previously indicated, but it does imply the absence evidence of selection to the employed or self-employed sectors.

Finally, to account for heterogeneity in individual unobservables, we add individual fixed effects to the first stage, which purges of individual-level time-invariant heterogeneity and confounding. The results are shown in Table 8. Our observation count drops to 478, which causes severely unstable estimates and a drastic change in results. The IMR is still significant for only employed workers but we cannot interpret this with certainty due to the small sample size.

7 Conclusions

This paper exploits two waves of individual-level panel data to examine the relationship between wages, British Nationality and ethnicities. We find that results corroborate with well-documented results the prior literature with regard to the relationship between wages and various demographic labels such as age, gender, location, etc. Self-employed workers tend to have lower and more volatile wages, but are not less educated, living in different areas, have a larger share of immigrants or working in particularly different industries. We find a negative relationship between cultural integration and self-employment, which is in line with findings from previous literature and theories. We find that the generation of immigrants has a positive relationship with wages only for the employed sector and for second generation immigrants only, which may reflect a number of factors such as language, ethnic enclave, parents becoming culturally integrated, such as immigrant success stories carried down to children, etc. Later generations of immigrants don't get a significant effect on wages, which may be due to the fact that they and their families are already fully assimilated. The self-selection effect to the employed or self-employed sectors is found to be significant and negative only for those in the employed sector but this is likely due to a small sample size issue for the self-employed sector.

A key assumption in our Heckman correction model is the assumption of normally distributed errors in both first and second stage models. This is potentiall a unreasonable assumption to make. In our future work, we aim to conduct a Box-Cox transformation on the dependent variables in the model to make them normally distributed. If effects were significantly different, we would have some evidence that the models' assumptions are inadequate and errors are indeed not normally distributed.

While the questionnaire asks a wide variety of interesting questions that other longitudinal surveys don't and a low attrition rate, the low response rates unfortunately either render these observations unusable, thus making inference unreliable if studies are ultimately carried out with a low observation count. We would like to have more conclusive data on English speaking ability, so that we could critically evaluate the language theory, and potentially get a deeper insight on the ethnic enclave theory by gaining access via special license. Improved data availability on variables related to measures of racism and race-induced xenophobia can act as potential causal channels for our topic of interest, which would make for a richer analysis.

8 Figures and Tables

Variable	Variable Code	Variable Value
Occupation	1	Argic./Forestry
	3	Energy/Water
	4	Mining
	7	Earth/Clay/Ston
	9	Mechanical Eng.
	10	Electrical Eng.
	11	Wood/Paper/Prit
	12	Clothing/Text.
	13	Food Industry
	14	Construction
	15	Constr. Relate
	16	Wholesale
	18	Retail
	19	Train System
	20	Communication/Entertainment
	21	Other Trans.
	22	Financial Inst
	23	Insurance
	24	Restaurants
	25	Service Indust
	27	Educ./Sport
	28	Health Service
	29	Legal Services
	30	Other Services
	31	Volunt./Church
	32	Priv. Household
	33	Public Admin.

Table 1: Variable Code for individual Occupations

Variable	Variable Code	Variable Value
Ethnicity	1	British/English/Scottish /Welsh/Northern Irish
	2	Irish
	4	Other White Background
	5	White and Black Caribbean
	6	White and Black African
	7	White and Asian
	8	Other Mixed Background
	9	Indian
	10	Pakistani
	11	Bangladeshi
	12	Chinese
	13	Other Asian Background
	14	Caribbean
	15	African
	16	Other Black Background
	17	Arab
	97	Other Ethnic Group

Table 2: Variable Code for individual Ethnicity

Variable	Variable Code	Variable Value
Highest Educational Attainment	1	Degree
	2	Other Higher Degree
	3	A-Level etc.
	4	GCSE etc.
	5	Other Qualification
	9	No Qualification

Table 3 (a): Variable Code for individual Education

$n = 227$	Emp. to Self-Emp.	Self-Emp. to Emp.
Immigrant	43	41
Native	66	77

Table 3 (b): Counts of Observations Changing Employed and Self-Employed Sectors

	Dependent variable:			
	1st Stage		2nd Stage	
	$\mathbb{P}(Z=0)$	$\mathbb{P}(Z=1)$	Y_{emp}	Y_{semp}
	(13)	(14)	(15)	(16)
(Intercept)	1.576***	−1.576***	7.341***	5.701***
	(0.511)	(0.511)	(0.175)	(2.083)
British Identity	0.030***	−0.030***		
	(0.009)	(0.009)		
First Job Self-Employed	−0.945***	0.945***		
	(0.109)	(0.109)		
Current Financial Situation	−0.066***	0.066***		
	(0.026)	(0.026)		
Future Financial Situation	0.109***	−0.109***		
	(0.028)	(0.028)		
Job Satisfaction	−0.081***	0.081***		
	(0.018)	(0.018)		
Number of Children	−0.066***	0.066***		
	(0.025)	(0.025)		
2nd Generation			0.037	0.284
			(0.024)	(0.354)
3rd Generation			0.047	0.423
			(0.044)	(0.640)
4th Generation			−0.019	0.754
			(0.036)	(0.459)
Healthy			0.074***	−0.063
			(0.018)	(0.254)
Part-Time Worker			−0.859***	−1.412***
			(0.020)	(0.283)
Female			−0.213***	0.066
			(0.017)	(0.286)
Age	−0.039**	0.039**	0.060***	−0.078
	(0.016)	(0.016)	(0.005)	(0.073)
Age^2	0.0001	−0.0001	−0.001***	0.001
	(0.0002)	(0.0002)	(0.0001)	(0.001)
Inverse Mills Ratio			−0.068	0.468
			(0.069)	(0.398)
Time Fixed Effects?	Yes	Yes	Yes	Yes
Education Qual. Dummies?	Yes	Yes	Yes	Yes
Ethnicity Dummies?	Yes	Yes	Yes	Yes
Job Industry Dummies?	Yes	Yes	Yes	Yes
Regional Dummies?	Yes	Yes	No	No
Parent Dummies?	Yes	Yes	No	No
Observations	5,567	5,567	4,824	743
Adjusted R^2			0.996	0.855
Log Likelihood	−1,715.138	−1,715.138		

Note: *p<0.1; **p<0.05; ***p<0.01

Table 4: Main regression

	Dependent variable:			
	1st Stage		2nd Stage	
	$\mathbb{P}(Z=0)$	$\mathbb{P}(Z=1)$	Y_{emp}	Y_{semp}
	(17)	(18)	(19)	(20)
(Intercept)	1.254***	−1.254***	7.294***	5.792***
	(0.395)	(0.395)	(0.159)	(1.976)
British Identity	0.025***	−0.025***		
	(0.007)	(0.007)		
Current Financial Situation	−0.026	0.026		
	(0.021)	(0.021)		
Future Financial Situation	0.072***	−0.072***		
	(0.023)	(0.023)		
Job Satisfaction	−0.060***	0.060***		
	(0.015)	(0.015)		
Number of Children	−0.094***	0.094***		
	(0.021)	(0.021)		
Born in UK			0.061***	0.173
			(0.019)	(0.268)
Healthy			0.076***	0.150
			(0.016)	(0.220)
Part-Time Worker			−0.860***	−1.058***
			(0.017)	(0.230)
Female			−0.212***	−0.080
			(0.015)	(0.238)
Age	−0.029**	0.029**	0.067***	−0.020
	(0.013)	(0.013)	(0.004)	(0.063)
Age^2	0.00003	−0.00003	−0.001***	0.0004
	(0.0001)	(0.0001)	(0.00005)	(0.001)
Inverse Mills Ratio			−0.227**	0.523
			(0.092)	(0.638)
Time Fixed Effects?	Yes	Yes	Yes	Yes
Education Qual. Dummies?	Yes	Yes	Yes	Yes
Ethnicity Dummies?	Yes	Yes	Yes	Yes
Job Industry Dummies?	Yes	Yes	Yes	Yes
Regional Dummies?	Yes	Yes	No	No
Observations	7,573	7,573	6,563	1,010
Adjusted R^2			0.996	0.859
Log Likelihood	−2,420.940	−2,420.940		
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 5: *ukborn* Regression Variation

	Dependent variable:			
	1st Stage		2nd Stage	
	$\mathbb{P}(Z=0)$	$\mathbb{P}(Z=1)$	Y_{emp}	Y_{semp}
	(9)	(10)	(11)	(12)
(Intercept)	1.590***	−1.590***	7.251***	5.086**
	(0.469)	(0.469)	(0.170)	(2.089)
British Identity	0.026***	−0.026***		
	(0.008)	(0.008)		
Current Financial Situation	−0.052**	0.052**		
	(0.024)	(0.024)		
Future Financial Situation	0.099***	−0.099***		
	(0.026)	(0.026)		
Job Satisfaction	−0.075***	0.075***		
	(0.017)	(0.017)		
Number of Children	−0.079***	0.079***		
	(0.024)	(0.024)		
2nd Generation			0.046**	0.232
			(0.023)	(0.325)
3rd Generation			−0.005	0.427
			(0.042)	(0.640)
4th Generation			−0.041	0.772*
			(0.035)	(0.451)
Healthy			0.065***	0.132
			(0.017)	(0.247)
Part-Time Worker			−0.851***	−1.204***
			(0.019)	(0.268)
Female			−0.219***	−0.038
			(0.016)	(0.270)
Age	−0.043***	0.043***	0.064***	−0.041
	(0.015)	(0.015)	(0.004)	(0.071)
Age^2	0.0001	−0.0001	−0.001***	0.001
	(0.0002)	(0.0002)	(0.0001)	(0.001)
Inverse Mills Ratio			0.058	0.523
			(0.082)	(0.523)
Time Fixed Effects?	Yes	Yes	Yes	Yes
Education Qual. Dummies?	Yes	Yes	Yes	Yes
Ethnicity Dummies?	Yes	Yes	Yes	Yes
Job Industry Dummies?	Yes	Yes	Yes	Yes
Regional Dummies?	Yes	Yes	No	No
Parent Dummies?	Yes	Yes	No	No
Observations	6,129	6,129	5,303	826
Adjusted R^2			0.996	0.853
Log Likelihood	−1,950.888	−1,950.888		
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 6: Robustness Check, Parent Education Regression

	Dependent variable:			
	1st Stage		2nd Stage	
	$\mathbb{P}(Z=0)$	$\mathbb{P}(Z=1)$	Y_{emp}	Y_{semp}
	(5)	(6)	(7)	(8)
(Intercept)	1.200***	−1.200***	7.470***	6.328***
	(0.438)	(0.438)	(0.170)	(1.897)
British Identity	0.031***	−0.031***		
	(0.008)	(0.008)		
First Job Self-Employed	−0.926***	0.926***		
	(0.099)	(0.099)		
Current Financial Situation	−0.041*	0.041*		
	(0.023)	(0.023)		
Future Financial Situation	0.085***	−0.085***		
	(0.025)	(0.025)		
Job Satisfaction	−0.058***	0.058***		
	(0.016)	(0.016)		
Number of Children	−0.096***	0.096***		
	(0.023)	(0.023)		
2nd Generation			0.055***	0.161
			(0.021)	(0.306)
3rd Generation			0.029	−0.106
			(0.040)	(0.583)
4th Generation			0.005	0.158
			(0.033)	(0.402)
Healthy			0.073	−0.090
			(0.016)	(0.229)
Part-Time Worker			−0.865***	−1.226***
			(0.018)	(0.248)
Female			−0.212	−0.014
			(0.016)	(0.259)
Age	−0.023	0.023	0.060***	−0.050
	(0.015)	(0.015)	(0.004)	(0.066)
Age^2	−0.00005	0.00005	−0.001	0.001
	(0.0002)	(0.0002)	(0.0001)	(0.001)
Inverse Mills Ratio			−0.331***	0.425**
			(0.073)	(0.421)
Time Fixed Effects?	Yes	Yes	Yes	Yes
Education Qual. Dummies?	Yes	Yes	Yes	Yes
Ethnicity Dummies?	Yes	Yes	Yes	Yes
Job Industry Dummies?	Yes	Yes	Yes	Yes
Regional Dummies?	Yes	Yes	No	No
Observations	6,633	6,633	5,747	886
Adjusted R^2			0.996	0.860
Log Likelihood	−2,086.004	−2,086.004		
Note:		*p<0.1; **p<0.05; ***p<0.01		

Table 7: Robustness Check, Self-Employed Dummy Regression Variation

	Dependent variable:			
	1st Stage		2nd Stage	
	$\mathbb{P}(Z=0)$	$\mathbb{P}(Z=1)$	Y_{emp}	Y_{semp}
	(21)	(22)	(23)	(24)
(Intercept)	17.328	−17.328	8.670***	7.812**
	(16.500)	(16.500)	(0.483)	(3.577)
British Identity	−0.064	0.064		
	(0.044)	(0.044)		
Current Financial Situation	−0.214*	0.214*		
	(0.114)	(0.114)		
Future Financial Situation	0.094	−0.094		
	(0.099)	(0.099)		
Job Satisfaction	−0.216***	0.216***		
	(0.064)	(0.064)		
Number of Children	−0.158	0.158		
	(0.152)	(0.152)		
2nd Generation			0.106	0.720
			(0.109)	(0.634)
3rd Generation			−0.187	1.181
			(0.221)	(1.234)
4th Generation			0.038	0.893
			(0.180)	(1.018)
Healthy			0.239**	−0.583
			(0.109)	(0.611)
Part-Time Worker			−1.008***	−1.006*
			(0.095)	(0.519)
Female			−0.219**	0.297
			(0.086)	(0.508)
Age	−0.384	0.384	0.005	−0.002
	(0.436)	(0.436)	(0.004)	(0.022)
Inverse Mills Ratio			−0.317*	0.719
			(0.167)	(1.003)
Time Fixed Effects?	Yes	Yes	Yes	Yes
Education Qual. Dummies?	Yes	Yes	Yes	Yes
Ethnicity Dummies?	Yes	Yes	Yes	Yes
Job Industry Dummies?	Yes	Yes	Yes	Yes
Regional Dummies?	Yes	Yes	No	No
Individual Dummies	Yes	Yes	No	No
Observations	478	478	254	224
Adjusted R^2			0.996	0.828
Log Likelihood	−273.071	−273.071		

Note: Age^2 omitted to avoid co-linearity issue — *p<0.1; **p<0.05; ***p<0.01

Table 8: Robustness Check, Individual Fixed Effects Regression

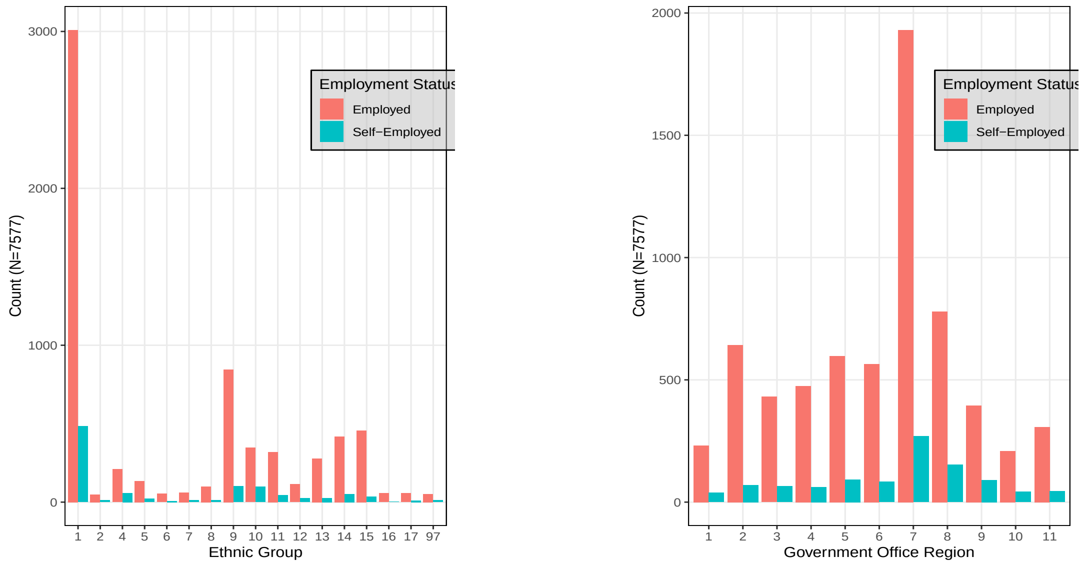

(a) Distribution of ethnic group (b) Distribution of Government Office Region

Figure (1) Key Variables Distributions

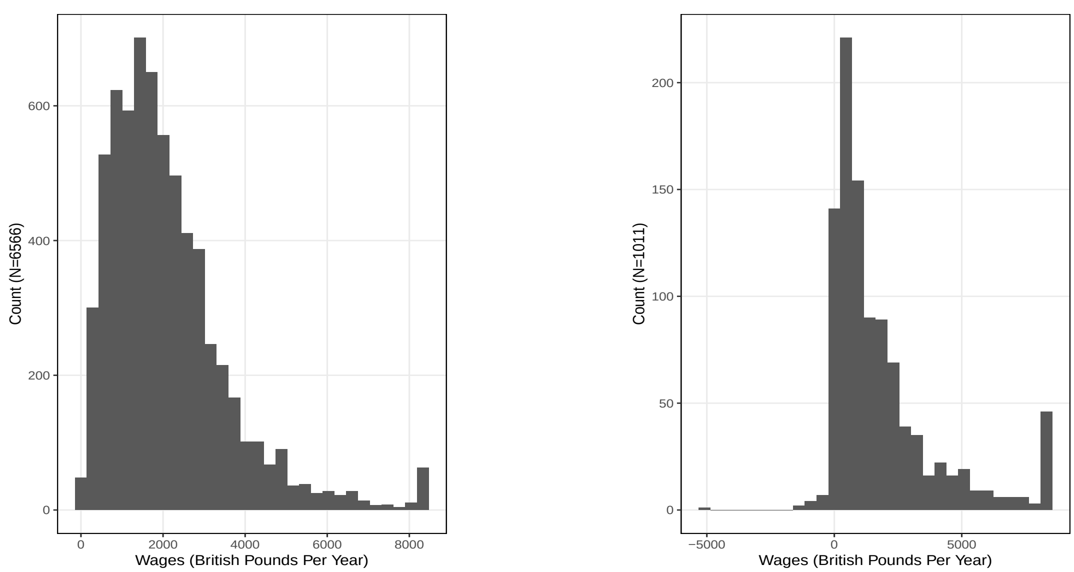

(a) Wage distribution of employed sector (b) Wage distribution of self-employed sector

Figure (2) Wage distributions

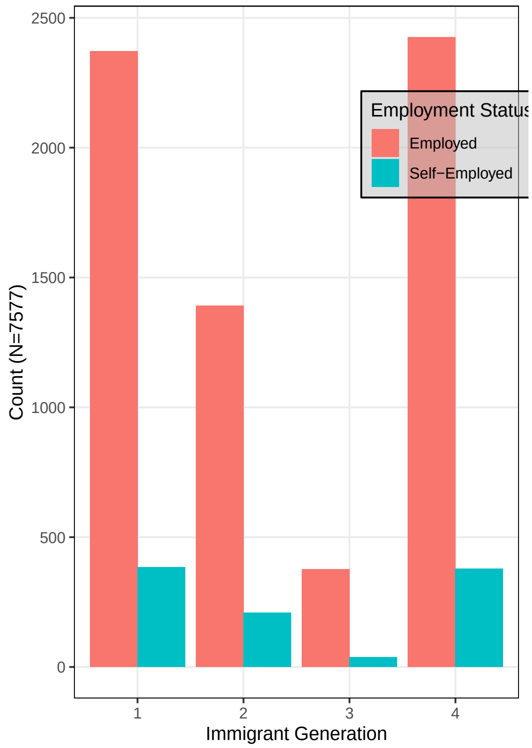

(a) Distribution of Immigrant Generation

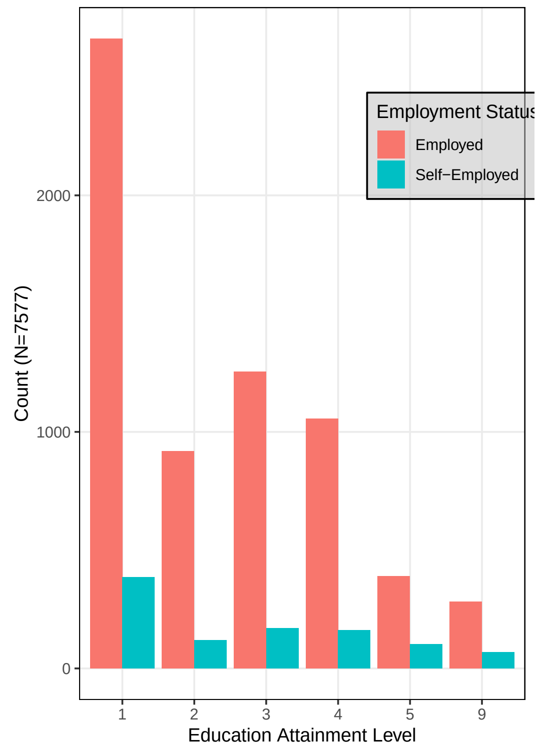

(b) Distribution of Highest Educational Attainment

Figure (3) Key Variable Distributions, Continued

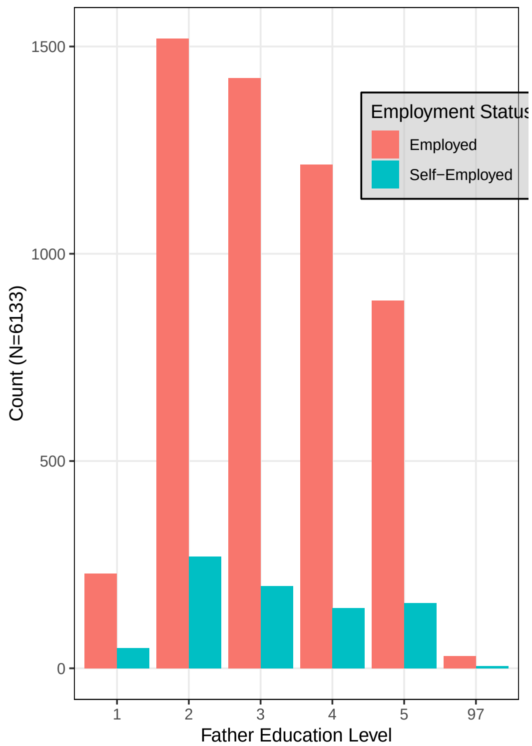

(a) Distribution of Father's Educational Attainment

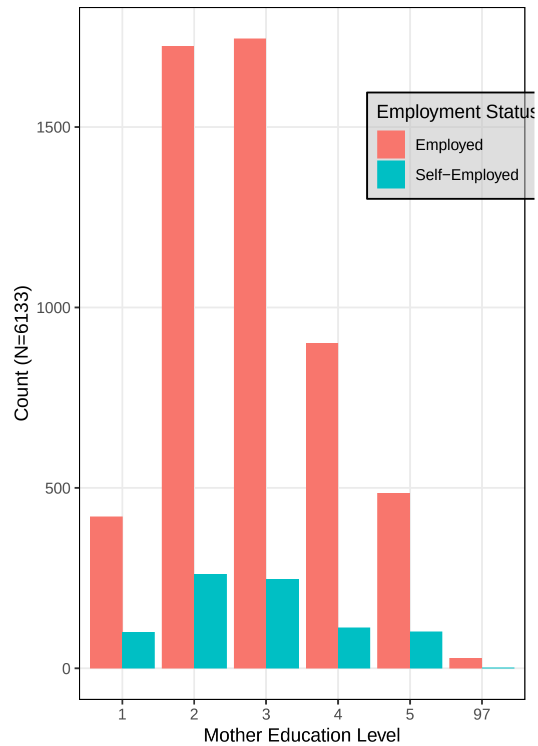

(b) Distribution of Mother's Educational Attainment

Figure (4) Parent's Educational Attainment

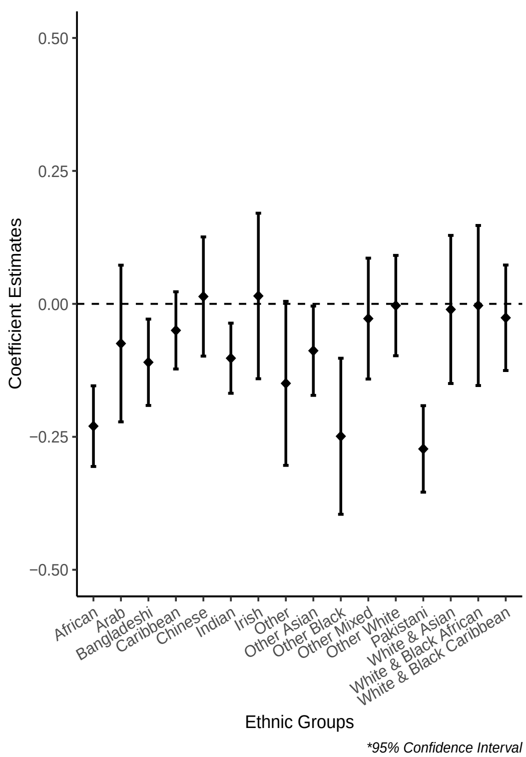

(a) Ethnic Group Effects on Wages of Employed

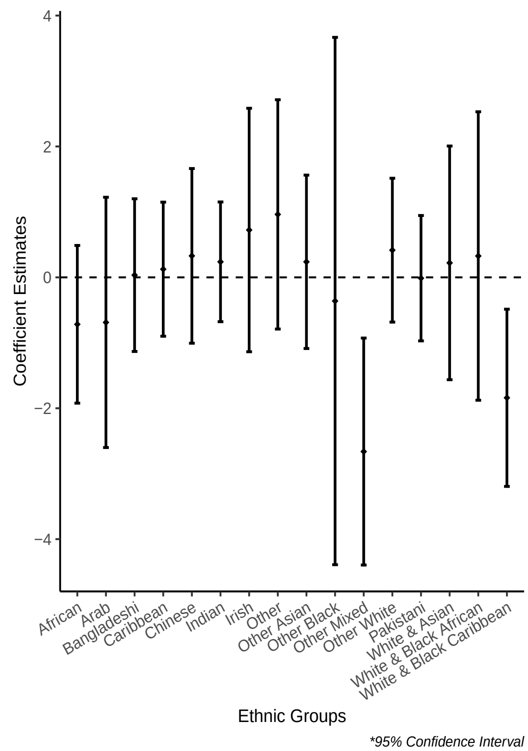

(b) Ethnic Group Effects on Wages of Self-Employed

Figure 5: Ethnic Group Effects on Wages (Relative to English/Scottish/Welsh/Northern Irish)

References

Aldrich, H., J. Cater, T. Jones, D. M. Evoy, and P. Velleman (1985): "Ethnic Residential Concentration and the Protected Market Hypothesis," *Social Forces*, 63, 996–1009.

Bartram, D. (2021): "Does the UK 'citizenship process' lead immigrants to reject British identity? A panel data analysis," *Ethnicities*, 21, 375–394.

Borjas, G. (1986): "The Self-Employment Experience of Immigrants," *Journal of Human Resources*, 21, 485–506.

Breunig, R., S. Hasan, and M. Salehin (2013): "The Immigrant Wage Gap and Assimilation in Australia: Does Unobserved Heterogeneity Matter?" *The Economic Record*, 89, 490–507.

Clark, K. and S. Drinkwater (2000): "Pushed Out or Pulled in? Self-Employment Among Ethnic Minorities in England and Wales," *Labour Economics*, 7, 603–628.

Clark, K., S. Drinkwater, and D. Leslie (1998): "Ethnicity and self-employment earnings in Britain 1973-95," *Applied Economics Letters*, 5, 631–634.

Cochran, T. C. (1972): "Ethnic Enterprise in America: Business and Welfare Among Chinese, Japanese, and Blacks. By Ivan H. Light. Berkeley, Cal., University of California Press, 1972. Pp. 209." *Business History Review*, 46, 376–377.

Evans, D. S. and L. S. Leighton (1989): "Some Empirical Aspects of Entrepreneurship," *The American Economic Review*, 79, 519–535.

Friedberg, R. M. (2000): "You Can't Take It with You? Immigrant Assimilation and the Portability of Human Capital," *Journal of Labor Economics*, 18, 221–251.

Hamilton, B. H. (2000): "Does Entrepreneurship Pay? An Empirical Analysis of the Returns to Self-Employment," *Journal of Political Economy*, 108, 604–631.

Heckman, J. J. (1979): "Sample Selection Bias as a Specification Error," *Econometrica*, 47, 153–161.

Johnson, N. L. (1949): "SYSTEMS OF FREQUENCY CURVES GENERATED BY METHODS OF TRANSLATION," *Biometrika*, 36, 149–176.

Nandi, A. and L. Platt (2015): "Patterns of minority and majority identification in a multicultural society," *Ethnic and Racial Studies*, 38, 2615–2634.

Portes, A. and M. Zhou (1993): "The New Second Generation: Segmented Assimilation and Its Variants," *The Annals of the American Academy of Political and Social Science*, 530, 74–96.

Wilson, K. L. and A. Portes (1980): "Immigrant Enclaves: An Analysis of the Labor Market Experiences of Cubans in Miami," *American Journal of Sociology*, 86, 295–319.

Labour Market Agglomeration Economies In The Context of Skill Heterogeneity: Evidence From England*

Stefanus Phan[1], Harish Vatsen[1], Sally Yang[1], and Marie Ogino[1]

[1]The London School of Economics and Political Science, London, UK

August, 2021

Abstract

Benefits from agglomeration economies are not uniform across a local labour market, but in fact exhibit significant heterogeneity along the dimension of occupational skills. We build on recent attempts to explore this notion by conducting a skill-based sub-sample analysis of the urban wage premium. We find that occupational skill profiles oriented towards analytical skills enjoy the highest urban wage premium, while those oriented towards physical and collaboration skills do not seem to enjoy any substantial urban wage premium.

*We thank the LSE SU Economics Society Research Division for the opportunity to write this paper, and for their unwavering support throughout the research process. In particular, we are deeply grateful to Weijian Zou and Mikael Johansson for their advice on both technical and administrative matters. We are thankful to Dr Guy Michaels for his valuable comments and LSE General Course student Jin Shuyun for her support in the preliminary stages of the project. We gratefully acknowledge the UK Data Service for data provision.

1 Introduction

Are there economic benefits to participating in thicker labour markets, such as those we see in big, populous cities? If so, who stands to gain the most? The advantages of labour market pooling, or the spatial concentration of workers and firms, have long been a topic of interest for researchers in urban economics. As early as 1890, Alfred Marshall advanced the notion that labour market pooling conferred economies of scale (see Marshall, 2009). These advantages, now known as agglomeration economies, have since been discussed and analysed in numerous theoretical and empirical works.

There are several channels through which agglomeration economies manifest (see Duranton and Puga, 2004). One can begin from Marshall's early emphasis that "a localized industry gains a great advantage from (offering) a constant market for skill". Both workers and employers benefit in this way: employers are less likely to face labour shortages, while workers are less likely to be unemployed. A related advantage is the impact of market size on worker-employer matching. Labour pooling makes searching less costly for both sides and enables more efficient matches on average. Another less obvious channel is the facilitation of sharing of skills and knowledge, which can be seen as human capital accumulation that makes workers more productive.

In this paper, we focus on what is perhaps the most prominent manifestation of agglomeration economies: the urban wage premium. Since the early 2000s, much scholarly attention has been paid to the question of why substantial wage differences exist between characteristically similar workers in urban and non-urban areas. This phenomenon, termed the urban wage premium, has been attributed to productivity advantages in urban areas due to agglomeration economies mentioned above. Its relative ease of measurability offers us an intuitive method for gauging the productivity effects of agglomeration economies. There are, however, alternative explanations of the urban wage premium. Wages in urban areas may be higher because of higher living costs, or that urban areas attract more productive workers to begin with (Yankow, 2006). The multi-faceted nature of wage determination makes the urban wage premium tricky to use as a proxy for urban productivity advantages. Nevertheless, urban economists have found ways to address these additional layers of complexity.

More recently, researchers have begun to explore the idea that labour markets, and the impact of agglomeration economies within a labour market, can be heterogeneous. In particular, labour markets can be extremely heterogeneous in terms of the types and levels of skills used by workers. It is plausible that firms and workers of different skill types and levels may benefit from agglomeration economies in different ways and to different extents. For instance, occupations with more non-routine skill profiles may entail a more complicated matching and screening process between employers and employees. We hypothesise that such occupations will benefit more from the search efficiencies generated by agglomeration economies. Building upon previous studies, we aim to derive more nuanced insights pertaining to the heterogeneity of labour market agglomeration economies across skill types.

Borrowing from the methodology of Choi (2020) we first delineate seven skill categories through factor analysis on data from the Skill and Employment Survey 2017, obtaining measures of each job type's intensity of use for each skill category. Combining these with data on job numbers from the Office for National Statistics (ONS), we construct an index of skill-based labour pooling. For each skill category, this index provides a measure of the extent of labour pooling within a county that is relevant to that skill category. For different skill categories, we conduct a panel regression analysis of the relationship between skill-based labour pooling and wages, using data from Understanding Society, the UK Household Longitudinal Study.

We make two main contributions. First, we apply a more nuanced categorisation of skills, enabling a more thorough breakdown of the nature of workplace skills and how they relate to agglomeration economies. Second, we introduce an alternative specification with interaction variables as an option for compactly analysing how the urban wage premium varies across different occupational skill profiles.

The rest of this paper is structured as follows. Section 2 provides an overview of the existing literature on labour market agglomeration economies and recent attempts to account for skill heterogeneity. Section 3 presents the data and the skill-based labour pool indices. Section 4 outlines our empirical methodology and main results. Section 5 explores robustness checks and extensions of the main analysis. Section 6 concludes.

2 Related Literature

Past literature provides a comprehensive discussion of agglomeration economies. Duranton and Puga (2004) outline a theoretical structure for characterising different microfoundations of urban agglomeration economies, including the facilitation and enhancement of matching as well as the generation, diffusion and accumulation of knowledge. Graham and Melo (2009) empirically evaluate the impact of agglomeration economies on hourly earnings using micro-level longitudinal worker data. Their findings reveal that greater market size and market potential are associated with increases in hourly wages, and that these economies of agglomeration decay with increasing radial distance from the worker's workplace. In an analysis of Italian Labour Force Survey micro-data, Di Addario (2011) obtains empirical findings that urbanisation and labour pooling raise job seekers' chances of finding employment. Bleakley and Lin (2012) find that workers in more densely populated areas, on average, earn higher wages and experience faster wage growth. Elvery (2010) finds that in bigger cities, the proportion of skilled labour rises instead of falling, despite higher labour costs, suggesting that skilled labour is more productive in bigger cities. All in all, extant literature generally lends credence to the idea that substantial agglomeration economies exist.

However, empirical evidence on agglomeration economies is not unanimous. De Blasio and Di Addario (2005), using Italian cluster mapping data, find that working within industrial clusters does not yield average wage premia for workers. Meanwhile, Haller and Heuermann (2020) investigate time spent unemployed after plant closures in Germany. Their results suggest that the negative effects of job competition exceed the benefits to job search of increased job density.

Of particular relevance to us are previous attempts to examine the urban wage premium. Glaeser and Mare (2001) use a wage regression with a dummy variable for cities with more than 5 million inhabitants, as well as interactions between this dummy and individual characteristics, and find that wages are 32% higher in large cities than in the hinterland. In a similar vein, Yankow (2006) uses a wage regression including a dummy variable for workers in large urban areas, and finds a 19% wage advantage. Other papers exploit the continuous variation in population density. Hirsch et al. (2020) estimate a 1.8-2.1% wage elasticity with respect to population density before controlling for local search frictions. Structural methods have also been used. Gould (2007) estimates a programming model which embeds residential location choices within a dynamic model of career decisions. Extant literature provides strong evidence that there are wage differentials between urban and non-urban areas.

However, few studies on thick labour markets have attempted to account for the heterogeneity of occupational skill profiles. Lazear (2009)'s introduction of a skill-weights approach set the stage for skill-based analysis of labour market phenomena. For instance, Gathmann and Schönberg (2010) calculate the distance in terms of skill profiles between occupations using combined German social security records career surveys. They find that task-specific human capital accounts for almost 52% of individual wage growth. In a similar vein, Geel et al. (2011) identify occupational clusters with similar skill combinations using data collected by the German Federal Institute for Vocational Education and Training. They find that occupational mobility within a cluster results in wage gains, while the reverse is true for across clusters. More recently, Choi (2020) employs a novel method of building indices of local labour pooling based on skill types, enabling an analysis of the heterogeneous impact of agglomeration economies across different skill types. Choi finds that the urban wage premium in South Korea is greatest for cognitive skill-oriented occupations and smallest for social skill-oriented occupations. Our paper primarily aims to extend Choi's analysis through the use of more specific skill categories.

Finally, we consider the true sources of the urban wage premium. While it appears a natural conclusion that greater market size causes productivity and hence wages to increase, D'Costa and Overman (2014) note the possibility of reverse causation: workers may be attracted to more productive areas, such that it is productivity that increases market size. Andersson et al. (2014) test the respective importance of the two sources and conclude that sorting is the main source of the urban wage premium. Combes et al. (2012) also find that individual skills account for a large fraction of existing spatial wage disparities.

3 Data and Derivation of Skill Variables

Before we perform a skill-based analysis of the urban wage premium, we first construct, for different categories of skills, measures of skill use intensity and indices of skill-based labour pooling. The former, which we call skill scores, are obtained through factor analysis on skill use data. The latter are obtained by combining these skill scores with information on job numbers in each county. This section details the data and methodology used to obtain both sets of variables.

3.1 Obtaining Skill Scores Through Factor Analysis

To gain an understanding of how different occupation groups vary in their skill use, we perform factor analysis on the skill use variables in the Skills and Employment Survey 2017 (SES2017). The SES2017 surveys a stratified random sample of 3306 workers aged 20-65 in Britain on their skills and employment. The data contains over 40 variables indicating the types of skills respondent's jobs involve and their intensity of use. These variables include reading and writing skills, physical skills such as dexterity and stamina, and interpersonal skills such as considering others' feelings and handling one's own feelings. Observations fall within a five-point Likert scale, with 5 representing the lowest intensity and 1 the highest. While informative, the large number of skill variables is likely to overcomplicate our econometric analysis. Factor analysis enables dimensionality reduction and a simpler categorisation of skills. This greatly improves the interpretability of our results in the subsequent econometric analysis.
We perform factor analysis using the principal-components method, which begins by assuming communalities to be 1. This method is preferred as it projects each data point onto only a specified number of principal components, thereby preserving as much of the variation in the skill use variables as possible. To discern the optimal number of factors, we make use of several tests, including Kaiser's criterion, Cattell's scree test and Horn's parallel analysis. These tests are not unanimous in their results, and our final set of options includes seven, nine and ten factors. With interpretability and ease of analysis being our key priorities, We choose to extract seven factors, the smallest number from these options. We provide a more detailed overview of the tests used in the appendix.

Additionally, we rotate the factors to improve their interpretability. We use oblique rotations as it is sensible to think that skill categories would be correlated with each other. For instance, those with greater use of literacy skills may tend to use problem solving skills more intensively in their jobs as well. This can be seen in Table 1 as well, where some correlations in the factor correlation matrix exceed 0.32. This suggests that there is significant overlap in variance among some of the factors, which warrants the use of oblique rotations (Tabachnick et al. 2007). We use the Promax rotation (with kappa set to 2) as the best choice for producing more interpretable factors.

Table 1: Factor Correlation Matrix

	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5	Factor 6	Factor 7
Factor 1	1						
Factor 2	.2908757	1					
Factor 3	.3894612	.2817047	1				
Factor 4	.3692299	.2764669	.3061636	1			
Factor 5	.1932531	.2022362	.2866553	.1202122	1		
Factor 6	.2326885	.1671465	.1657101	.1955546	.1282455	1	
Factor 7	-.1787539	-.1315919	.0168925	-.0635449	-.0315993	.0353991	1

A correlation matrix for the 7 factors (as yet unnamed in this stage of the factor analysis)

Using the factor loadings in Table 2 as a guide, the factors can be interpreted as distinct skill categories. Factor 1 comprises communication skills, such as reading and writing. Factor 2 represents managerial skills, such as motivating staff and planning the activities of others. Factor 3 represents analytical skills, such as spotting faults and analysing complex problems. Factor 4 represents skills which involve interaction with clients, such as persuading others and looking the part. We henceforth refer to Factor 4 as "client-facing skills" for short. Factor 5 comprises quantitative skills, such as using a computer and advanced mathematics. Factor 6 includes collaboration skills, such as teamwork and cooperation. Finally, Factor 7 represents physical skills, such as stamina and strength. A more detailed breakdown of how we interpret the seven factors is provided in the appendix.

Of note is how our factors compare to those obtained by Choi. Client-facing skills and collaboration skills can be seen as subgroups of Choi's social skills category. Communication skills, analytical skills and quantitative skills correspond most closely to Choi's cognitive skills category. Managerial skills do not correspond well to any of Choi's skill categories. This is attributable to differences in the design of survey questions.

We extract these seven factors as standard normal variables each representing a skill category, as described above. These variables can be treated as skill scores reflecting the intensity of skill use in the corresponding category. A score of 0 reflects an average use of skills in the corresponding category, and a positive (negative) score would reflect an above-average (below-average) intensity of use. For each skill category, an average score is calculated across each occupational sub-major group, as defined under the 2010 version of the United Kingdom's Standard Occupational Classification (SOC) system at the 2-digit level. Table 3 contains a breakdown of these mean skill scores.

Table 2: Rotated Factor Loadings

	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5	Factor 6	Factor 7
cpeople				.6174687			
cteach		.3021074					
cspeech							
cpersuad				.3940895			
cselling	-.3558124			.5521052	.4491031		
ccaring				.6490513			
cteamwk						.7872843	
clisten						.7667144	
cstrengt							.8727746
cstamina							.874341
chands							.7921925
ctools							.73421
cproduct			.3893178		.3270256		
cspecial			.456964				
corgwork							
cusepc					.3490464		-.4169724
cfaults			.7985438				
ccause			.8402508				
csolutn			.8102671				
canalyse			.5537609				
cplanme			.4202342			-.3339621	
cplanoth		.4672896					
cmytime			.436499				
cahead	.3325107		.3780797				
cread	.711688						
cshort	.7428141						
clong	.7419553						
cwrite	.7757262						
cwritesh	.7304396						
cwritelg	.6820633						
ccalca					.7320811		
cpercent					.8002813		
cstats					.7281997		
ccoop						.7395772	
cmotivat		.9646276					
cthings		.9539752					
ccoach		.9699866					
ccareers		.9451238					
cfuture		.8817072					
cmefeel				.5714723			
cothfeel				.5889769			
clookprt				.7068302			
csoundprt				.7240926			

Loadings with absolute value less than 0.3 are omitted for ease of reading

Table 3: Mean Skill Scores for Occupational Submajor Groups

2-Digit SOC	Communication	Managerial	Analytical	Client-Facing	Quantitative	Collaboration	Physical
11	0.240	1.207	0.397	0.454	0.579	0.125	-0.432
12	0.0643	0.828	0.314	0.385	0.401	-0.230	0.0241
21	0.209	0.290	0.694	-0.410	0.693	0.0542	-0.762
22	0.648	0.481	0.389	0.628	-0.0185	0.503	0.214
23	0.878	0.495	0.131	0.715	0.134	0.0731	-0.483
24	0.744	0.283	0.368	0.170	0.125	0.0843	-1.084
31	0.0545	-0.129	0.555	-0.501	0.448	0.237	-0.353
32	0.436	-0.0294	0.0819	0.622	-0.331	0.125	-0.334
33	0.595	-0.0926	-0.118	0.333	-0.731	0.606	0.247
34	-0.0246	-0.270	0.444	0.265	-0.201	-0.715	-0.0849
35	0.217	0.201	0.255	0.328	0.476	-0.00661	-0.832
41	0.181	-0.316	-0.0978	-0.301	0.204	0.210	-0.728
42	0.103	-0.309	-0.309	0.229	-0.146	0.298	-0.554
51	-0.270	-0.187	0.339	-0.387	0.139	-1.139	1.156
52	-0.0508	-0.256	0.608	-0.595	0.208	-0.0183	0.917
53	-0.349	-0.396	0.310	-0.252	0.0690	-0.715	1.132
54	-0.384	0.208	0.143	-0.140	-0.0878	-0.152	0.949
61	0.360	-0.268	-0.229	0.299	-0.685	0.263	0.349
62	-0.659	-0.221	-0.338	0.235	-0.462	-0.319	0.414
71	-0.923	-0.239	-0.615	0.161	0.256	0.154	0.235
72	-0.110	-0.423	0.0285	0.147	0.240	0.563	-0.562
81	-0.337	-0.0908	0.103	-0.967	-0.0889	0.0830	0.832
82	-0.450	-0.607	-0.683	-0.458	-0.547	-0.509	0.528
91	-0.819	-0.550	-0.768	-1.172	-0.440	-0.209	0.910
92	-0.948	-0.479	-0.918	-0.604	-0.636	-0.211	0.626

The sub-major group corresponding to each 2-Digit SOC code can be found on this ONS webpage

3.2 Deriving Skill-Based Indices of Local Labour Pooling

To obtain skill-based measures of local labour pooling within areas, we combine our extracted factors with statistics derived from the Annual Population Survey by the UK Office for National Statistics. Specifically, we use estimates of the number of jobs in an area for each occupational sub-major group.

Taking each district to be an area would allow for more granular analysis. However, district-level analysis is less robust to the possibility of workers working in one district and living in another. Since our access to Special License UKHLS data only provides information on where one lives in and not where one works, we inevitably face an issue of some individuals being assigned the wrong area, i.e. the area they reside in rather than the area they work in. This would cause measurement error in population (density) and the skill-based labour pooling indices, which are area-dependent variables. Using districts as areas would compound this issue, compared to using broader subdivisions, such as counties or regions.

It may also be unrealistic to treat local labour markets as being confined to singular districts. A case in point is London. Each of its boroughs is considered a district, but it is very common for one to live in one borough and commute to another for work.

Given the above considerations, we calculate our skill-based labour pooling indices at the county level. While we do still face the possibility of one commuting to another county to work, it is now a much smaller concern compared to at the district level, given that counties are a broader subdivision of England. ONS statistics on commuting patterns offers some affirmation of this intuition. Another complication is the existence of unitary authorities, which are treated as both districts and counties. However, they are generally a small part of a larger ceremonial county. Therefore, we group all districts (including unitary authorities) into their respective ceremonial counties. There are 48 ceremonial counties in total, including six metropolitan counties and 42 non-metropolitan counties. One difference in our classification of counties compared to the official classification of ceremonial counties is in our treatment of London. While the official classification separates London into the City of London and Greater London, we separate London into Inner and Outer London. This is to avoid complications from the City of London's extremely small area and population size, while still broadly acknowledging the differences in labour market thickness within London.

Using this data along with the skill scores calculated prior, we modify Choi's method to derive a skill-based index of local labour pooling:

$$I_{jst} = \frac{\sum_l E_{jlt} \times K_{ls}}{A_j} \tag{1}$$

where I_{jst} is the skill-based labour pool index for skill category s in county j at time t, E_{jlt} is the number of workers in occupation group l and county j at time t, K_{ls} is the skill score in skill category s for occupational sub-major group l and A_j is the area of county j (which we assume to be constant throughout the analysis time frame).

The numerator is a sum of employment numbers for each occupational sub-major group, weighted by their skill score. Our modification is to have this weighted sum be divided by county area. The entire expression can then be interpreted as a measure of job density relevant to the skill category concerned. For this calculation only, we treat all negative skill scores as 0, so that we do not subtract from the local labour pooling measure when accounting for occupation groups with below-average intensity of skill use. In sum, for each skill category, the index measures the density of workplace skill use relevant to that skill category in a county.

In Table 4, we provide an overview of the seven skill-based labour pool index values for each ceremonial county for the most recent year. Unsurprisingly, Inner and Outer London have some of the largest values for all seven skill categories. Similarly, all six of the metropolitan counties rank near the top for each skill category. That said, the "metropolitan vs non-metropolitan" distinction can be misleading, as non-metropolitan counties can also have urban areas within them. For instance, Surrey and Nottinghamshire are non-metropolitan counties with a large portion of urban areas, and correspondingly high scores on the skill-based index.

Table 4: Skill-Based Labour Pooling Indices For Ceremonial Counties

county	Communication	Managerial	Analytical	Client-Facing	Quantitative	Collaboration	Physical
Bedfordshire	59.83	61.05	57.81	57.76	56.17	32.23	70.84
Berkshire	89.38	95.93	95.41	78.37	97.64	41.80	78.79
Bristol	608.59	512.25	541.11	501.41	476.14	277.59	505.77
Buckinghamshire	51.23	63.31	51.81	49.64	52.86	26.80	47.41
Cambridgeshire	26.71	28.86	28.68	22.93	28.63	13.24	31.45
Cheshire	50.30	55.65	50.62	47.52	51.07	27.65	55.14
Cornwall	14.22	13.88	13.95	14.64	12.55	7.70	25.81
County Durham	32.39	28.65	29.66	31.00	27.51	20.38	48.78
Cumbria	6.17	6.53	6.74	5.72	5.92	3.23	12.61
Derbyshire	37.59	38.73	38.38	35.98	36.66	21.06	55.53
Devon	17.65	17.57	15.80	16.90	14.94	9.33	24.35
Dorset	31.58	34.55	31.14	30.99	28.64	16.81	39.19
East Riding of Yorkshire	23.03	22.99	21.07	22.46	20.14	14.30	36.88
East Sussex	56.47	58.34	46.40	55.69	47.68	25.72	48.21
Essex	51.25	57.44	49.78	48.65	49.80	27.94	60.71
Gloucestershire	32.95	34.16	32.57	30.73	30.55	17.20	38.96
Greater Manchester	233.04	217.55	210.51	218.45	202.79	131.89	281.96
Hampshire	53.23	56.91	51.67	49.48	52.59	27.33	55.97
Herefordshire, County of	9.04	9.78	9.52	9.42	8.74	5.36	14.79
Hertfordshire	86.83	98.06	81.56	79.85	80.47	35.43	68.35
Inner London	1,608.72	1,742.46	1,552.19	1,430.78	1,437.01	619.98	956.34
Isle of Wight	32.48	30.55	30.53	33.41	27.53	18.37	50.04
Kent	51.61	54.90	44.88	48.48	46.07	24.71	53.88
Lancashire	43.55	43.69	38.01	42.36	39.96	27.21	61.82
Leicestershire	51.62	53.86	47.22	49.27	46.44	26.79	65.00
Lincolnshire	12.06	13.27	11.91	12.29	11.46	7.12	23.09
Merseyside	223.81	203.69	194.68	212.12	187.74	136.92	286.68
Norfolk	14.08	15.72	14.24	14.34	14.55	8.78	24.95
North Yorkshire	12.54	12.15	10.57	11.74	10.23	6.79	15.62
Northamptonshire	29.88	30.22	27.47	26.40	28.72	17.78	47.49
Northumberland	6.24	6.46	6.22	5.96	5.45	3.51	8.88
Nottinghamshire	52.85	51.11	49.02	50.40	45.19	29.24	74.73
Outer London	504.64	537.43	479.05	460.70	476.91	235.69	449.74
Oxfordshire	33.56	37.05	35.77	29.18	36.77	15.51	26.24
Rutland	8.68	9.03	8.91	8.49	9.10	3.82	11.13
Shropshire	15.27	16.77	14.70	14.80	14.15	7.84	21.25
Somerset	24.43	25.95	22.70	23.47	22.68	12.85	32.05
South Yorkshire	92.67	86.37	81.41	90.36	78.43	58.41	122.64
Staffordshire	41.82	45.37	40.87	38.95	39.50	21.76	54.39
Suffolk	20.77	20.98	19.03	19.42	17.66	10.64	27.65
Surrey	91.23	109.13	92.41	82.84	92.98	34.48	51.45
Tyne and Wear	204.49	185.65	190.02	196.31	183.89	132.36	268.38
Warwickshire	29.39	33.15	31.56	26.81	32.55	14.67	34.38
West Midlands	300.08	260.09	276.00	272.93	251.93	179.08	448.10
West Sussex	45.36	50.23	44.11	44.80	44.79	22.57	47.40
West Yorkshire	112.69	104.88	106.42	108.23	103.40	64.34	153.49
Wiltshire	21.32	24.15	22.43	21.65	21.85	11.71	28.80
Worcestershire	35.23	36.21	33.54	31.87	32.17	17.58	43.73

By themselves, the scores on the skill-based labour pooling index tell us little about regional skill specialisations, as they strongly correlate with county population. They are more reflective of absolute advantages in skill stock, as opposed to comparative advantages. Instead, we follow Choi's method of dividing a county's share of the England-wide skill-weighted sum of jobs (the numerator in our formula for the skill-based index) by that county's share of England's total population. Comparing this ratio across counties, we do find distinct patterns of regional specialisation by workplace skills. Figures 4 to 10 in the appendix (7.4) visualise these specialisations for each skill category using filling maps. Most prominently, Surrey has a relatively high degree of specialisation in analytical, quantitative and management skills, but one of the lowest in physical skills. This aligns with Surrey's status as a financial and management hub of the UK. Another illustrative example is East Sussex which has some of the highest degrees of specialisation in client-facing skills. This echoes the importance of tourism to their local economy. A final example is that of Cumbria, which specialises the most in physical skills. This can be explained by the predominance of industry in Cumbria's local economy.

3.3 Panel Data for Regression Analysis

For our analysis of the urban wage premium, we use data from the Understanding Society, the UK Household Longitudinal Study (UKHLS). The UKHLS is the successor to the British Household Panel Survey (BHPS). Since 2009, it has involved yearly interviews of 40000 households, including 8000 households continuing from the BHPS. We use waves 3 to 10 (or 2012 to 2019) of the UKHLS for our study. This is done for two reasons. Firstly, the Standard Occupational Classification coding system changes from the 2000 version to the 2010 version from wave 2 to wave 3. Secondly, keeping the analysis timeframe to the years 2012 to 2019 fits better with the variables we derived from factor analysis use 2017 data.

The UKHLS provides rich information on individual, household and occupational level characteristics. Using the UKHLS, we divide the measure of average wages earned per week by average hours worked per week to obtain a measure of hourly wage. In this calculation, we account for inflation using the Consumer Price Index, setting July 2015 as our reference level, and for any self-employed labour activity reported.

We assign each individual a set of seven skill-based labour pooling indices by mapping each the individual's district of residence. This assumes that individuals work in the same ceremonial county they live in. As discussed in subsection 3.2, we believe that setting our identifying variation at the county level suitably minimises the possibility of violating this assumption.

4 Empirical Methodology and Results

4.1 Analysis of the Urban Wage Premium Without Allowing For Skill Heterogeneity

We first test the urban wage premium hypothesis without allowing for skill heterogeneity. This involves examining the relationship between wage and population, as well as between wage and population density. The regression model used can be expressed as:

$$\ln w_{ijt} = \gamma \ln P_{jt} + \beta X_{ijt} + \delta_t + \epsilon_{ijt} \tag{2}$$

Above, w_{ijt} is the wage of individual i in county j at time t. Depending on the analysis, P_{jt} represents population or population density of a county j at time t. X_{ijt} is a vector of controls for individual and occupational characteristics, such as age, education, gender and industry. A squared age term is added to allow for the effect of age to change non-linearly as the individual grows older, in line with the Mincer earnings function. We also consider controlling for tenure, but leave it out in our main analyses, as tenure is indirectly derived from other variables in the dataset, and is missing for a large fraction of individuals.

δ_t represents time fixed effects, such as macro shocks or underlying universal trends in wage. There is also the option to include individual fixed effects μ_i, which would control for unobserved time-invariant characteristics such as innate ability or personality traits. However, a large fraction of surveyed individuals did not change their county of residence within the sample timeframe. Adding individual fixed effects may eliminate most of the identifying variation in our data, causing the variance introduced to outweigh the possible reduction of bias. Thus, we first follow Choi's decision to leave out individual fixed effects. That said, we do check the soundness of this decision in section 5.5.

As seen in Table 5, regressing log of wage on log of population and on log of population density respectively gives results which are consistent with previous literature. We find an elasticity of 1.86% for wage with respect to population, and an elasticity of 0.823% for wage with respect to population density. As a benchmark, we note that Choi's estimates are 3.2% and 2.2% respectively. While the elasticity with respect to population density being markedly lower is consistent with Choi's results, Choi does obtain higher estimates than us for both elasticities. This difference in estimates may be due to differences in institutional or labour market factors between England and South Korea.

Table 5: Regression Results for Population and Population Density

	(1) Log Real Wage	(2) Log Real Wage
Log Population	0.0186***	
	(0.00373)	
Log Population density		0.00823***
		(0.00209)
Age	0.0465***	0.0465***
	(0.00123)	(0.00123)
Age^2	-0.000470***	-0.000471***
	(0.0000140)	(0.0000140)
Number of Children	-0.0218***	-0.0218***
	(0.00289)	(0.00289)
N	143905	143905

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

4.2 Skill-Based Subsample Analysis of the Urban Wage Premium

We next perform a skill-based subsample analysis of the urban wage premium hypothesis using the skill score and skill-based labour pooling variables made earlier. When estimating the urban wage premium associated with skill category s, we use a subsample of individuals whose skill score for that skill category is greater than or equal to zero. In other words, our subsample includes all individuals with at least average use of skill category s in their occupation.

The regression model used for a skill category s can be expressed as:

$$\ln w_{ijt} = \gamma lnI_{sjt} + \theta K_{sit} + \beta X_{ijt} + \delta_t + \epsilon_{ijt} \tag{3}$$

Where P_{jt} is replaced by I_{sjt} as the main regressor, and K_{sit}, the mean skill score for the 3-digit SOC code of an individual i's job at time t, is included as an additional control. We use the mean skill score corresponding to the 3-digit SOC code as it allows us to treat occupation types as granularly as possible while maintaining the reliability of the estimated mean skill score. Using 4-digit SOC codes would have resulted in intolerably unreliable estimates of mean skill scores as the sample used from the SES 2017 would not have been large enough to provide a sufficient number of data points for each SOC code.

In this model, γ can be interpreted as an elasticity of wage with respect to the labour pooling index for skill category s. Comparing estimates of this wage-index elasticity across the different subsample regressions, provided in Tables 6 to 8, sheds some light into how the urban wage premium varies across occupational skill profiles. The highest estimates of the wage-index elasticity we obtain are for analytical skills and communication skills (in Table 6, at 1.8% and 1.4% respectively. This suggests the urban wage premiums associated with these skill categories are strongest. We also find a strong urban wage premium associated with managerial skills, client-facing skills and quantitative skills, obtaining elasticity estimates of 1.14%, 1.22% and 0.96% respectively (in Table 7). In contrast, our estimates of the wage-index elasticity for physical skills and collaboration skills (in Table 8) are not statistically significant at even the 10% significance level. This suggests that an urban wage premium associated with physical skills and collaboration skills is virtually non-existent.

Our results point to a substantial divergence in the urban wage premium across occupational skill profiles, which in turn points to a stark heterogeneity in the benefits received from agglomeration economies. The main pattern we identify in this divergence is that more non-routine skills, such as analytical and communication skills, are associated with a higher urban wage premium than more routine or generic skills, such as physical and collaboration skills. This pattern can be rationalised using our understanding of the channels through which agglomeration economies lead to productivity advantages. For instance, occupations that require a non-routine skill profile may entail a more complicated matching process. Such occupations would therefore benefit more from the search efficiencies generated by agglomeration economies.

Also of interest are our estimated coefficients on the mean skill score variables. We find positive and statistically significant coefficients for communication, client-facing, quantitative and managerial skills, at 0.376, 0.323, 0.257 and 0.114 respectively. For communication skills, the coefficient of 0.376 can be interpreted as follows: for occupations with at least average use of communication skills, an increase in one's communication skill score by one standard deviation increases one's wage by 37.6%, all else equal. Meanwhile, we do not find a significant coefficient for analytical skills. When considered alongside the fact that the estimated wage-index elasticity was highest for analytical skills, this suggests agglomeration economies are especially important for analytical occupational skill profiles. Finally, we find negative and statistically significant coefficients for collaboration and physical skills, at -0.228 and -0.253 respectively. As an illustration, for physical skills, this can be interpreted as a 25.3% fall in wage associated with a one standard deviation increase in one's physical score, for occupations with at least average use of physical skills to begin with. The results obtained are similar whether or not we control for industry. This addresses any concerns that our results could be driven largely by inter-industry wage differentials.

The results obtained here also seem to affirm our decision to use more granular skill categories. While our results generally agree with Choi's, our breakdown of social skill categories into client-facing and collaboration skills offers further nuance to his findings regarding social skills. While Choi's results suggest that benefits from local labour pooling are non-existent for social skills, we find that the urban wage premium is non-existent for collaboration skills but relatively high for client-facing skills, which are less routine.

At this point it is prudent to keep in mind some of the limitations of our analysis. For one, as factors such as living costs and innate ability are not explicitly controlled for, they are highly plausible sources of omitted variable bias. One can also expect measurement error in our index variables, which can generate attenuation bias in our elasticity estimates. Meanwhile, measurement error in the wage variable would reduce the precision of our estimates. We attempt to address some of these limitations in the following section.

Table 6: Regression Results for Analytical and Communication Skills

	(1) Log Real Wage	(2) Log Real Wage
Log Index of Analytical Skills *(County Level)*	0.0181***	
	(0.00430)	
Mean Analytical Skill Score *(3 digit SOC)*	0.0368	
	(0.0285)	
Log Index of Communication Skills *(County Level)*		0.0140***
		(0.00375)
Mean Communication Skill Score *(3 digit SOC)*		0.376***
		(0.0194)
Age	0.0463***	0.0471***
	(0.00290)	(0.00262)
Age^2	-0.000527***	-0.000510***
	(0.0000333)	(0.0000305)
Number of Children	-0.0351***	-0.0250***
	(0.00631)	(0.00560)
N	44577	42745

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Columns 1 and 2 pertain to analytical and communication skills respectively

Table 7: Regression Results for Managerial, Client-Facing and Quantitative Skills

	(1) Log Real Wage	(2) Log Real Wage	(3) Log Real Wage
Log Index of Managerial Skills	0.0112*		
	(0.00470)		
Mean Managerial Skill Score	0.114***		
	(0.0158)		
Log Index of Client-Facing Skills		0.0122***	
		(0.00352)	
Mean Client-Facing Skill Score		0.323***	
		(0.0207)	
Log Index of Quantitative Skills			0.00962*
			(0.00414)
Mean Quantitative Skill Score			0.257***
			(0.0214)
Age	0.0473***	0.0294***	0.0406***
	(0.00344)	(0.00225)	(0.00271)
Age^2	-0.000531***	-0.000320***	-0.000467***
	(0.0000395)	(0.0000266)	(0.0000317)
Number of Children	-0.0307***	-0.0227***	-0.0222***
	(0.00693)	(0.00534)	(0.00620)
N	33968	50275	44023

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Columns 1, 2 and 3 pertain to managerial, client-facing and quantitative skills respectively

Table 8: Regression Results for Physical and Collaboration Skills

	(1) Log Real Wage	(2) Log Real Wage
Log Index of Physical Skills	-0.000573	
	(0.00429)	
Mean Physical Skill Score	-0.253***	
	(0.0171)	
Log Index of Collaboration Skills		0.00444
		(0.00323)
Mean Collaboration Skill Score		-0.228***
		(0.0251)
Age	0.0217***	0.0372***
	(0.00226)	(0.00204)
Age^2	-0.000254***	-0.000407***
	(0.0000268)	(0.0000242)
Number of Children	-0.0171**	-0.0168***
	(0.00556)	(0.00471)
N	41114	57901

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Columns 1 and 2 pertain to physical and collaboration skills respectively

5 Robustness Checks and Extensions

5.1 Clustering Standard Errors

It is highly plausible for wages to be positively correlated over time at the individual level. A naive analysis will systematically underestimate standard errors, thereby giving a false boost to the precision of our estimates. We correct for this by clustering standard errors at the individual level. The resultant increase in standard errors does have important implications for some of our results. Now, our estimate of wage-index elasticity for managerial skills is not significant at the 10% level. For physical skills, the p-value increases to over 0.8, which further reduces our confidence that there is any urban wage premium associated with physical skills at all.

5.2 Standard Occupational Classification Codes: 2-digit vs 3-digit

We assume 3-digit SOC codes classify occupations more accurately than 2-digit SOC codes. Therefore, our regression uses mean skill scores for 3-digit SOC codes. We check this assumption by replacing the mean skill score variables for 3-digit SOC codes with their 2-digit counterparts. This has varying effects on the estimated wage-index elasticities across our skill-based regressions, though the differences are small in magnitude. For instance, the estimated wage-index elasticity rises from 0.96% to 1.4% for quantitative skills, but falls from 1.8% to 1.77% for analytical skills. Perhaps a vote of confidence for using 3-digit SOC codes is the observation that the R-squared and adjusted R-squared is generally lower when using mean skill scores for the 2-digit SOC codes, suggesting a worse model fit.

5.3 Using Alternative Formulations of Real Hourly Wage

As mentioned in subsection 3.3, we derived measures of real wage by dividing average weekly earnings by average weekly hours worked, and then adjusting these estimates using the CPI index. We test the robustness of our results by using several alternative formulations of real hourly wage, such multiplying overtime hours worked by 1.5 to account for the possibility of higher overtime pay. We also choose different time points to be the reference price level when accounting for inflation. In all of these tests, we find no significant change in our results.

5.4 Checking Sensitivity to Skill Score Threshold for Subsample Analysis

When constructing our subsamples for the skill-based analyses of the urban wage premium, we use 0 as a lower bound for the relevant skill score of individuals in the subsample. We rerun each regression using several alternative values for the threshold. We find that the estimated wage-index elasticity generally changes in sensible directions when we adjust the threshold. For instance, increasing the threshold to 0.2 increases the estimated wage-index elasticity while decreasing it to -0.2 decreases it. This is in line with the idea that occupations with more intensive use of a particular skill category would also enjoy a greater urban wage premium associated with that skill category.

5.5 Individual Fixed Effects

As mentioned in section 4.1, we have reason to believe, *a priori*, that incorporating individual fixed effects, μ_i, into our regression analysis will be inappropriate. It is nevertheless useful to examine how our results change when adding μ_i, since it has the advantage of controlling for relevant time-invariant factors such as ability. Indeed, when including individual fixed effects, the coefficient on the main regressor is rendered insignificant across all models.

This change in results may be regarded as consistent with previous literature that has tested competing explanations of the urban wage premium. Andersson et al. (2014) find a drastic fall in their estimated wage-density elasticities when including individual fixed effects, echoing the idea that spatial sorting is also an important cause of the urban wage premium. However, the change in our results can also be attributed to a loss of identifying variation since a large fraction of individuals in our sample do not change county within the sample timeframe. Attenuation bias, generated by mismeasurement in our index variables, is another possible explanation.

5.6 Accounting for Correlations between Skill Categories

We attempt an alternative regression analysis method for investigating how the urban wage premium varies across occupational skill profiles. In contrast with the main analysis, this method simply uses population density instead of the skill-based labour pooling indices. However, it incorporates mean skill score variables for all seven skill categories in one regression, as well as an interaction variable between population density and each mean skill score. Compared to the main analysis, this gives it the added advantage of allowing for correlations between skill categories. This method does not use the skill-based labour pooling indices, and hence discards any additional information these variables provide. It also does not address the sources of omitted variable bias faced by the main analysis. However, we believe it is still useful as an extension as it enables a compact picture of how the urban wage premium is influenced jointly by the use of skills from all seven skill categories.

The regression equation described above can be written as follows:

$$\ln w_{ijt} = \gamma lnP_{jt} + \sum_{s} \theta K_{sit} + \sum_{s} \omega_s K_{sit} lnP_{jt} + \beta X_{ijt} + \delta_t + \epsilon_{ijt} \tag{4}$$

Here we denote the set of mean skill scores as the summation $\sum_s \theta K_{sit}$ and the set of interactions between skill score and population density as the summation $\sum_s \omega_s K_{sit} lnP_{jt}$. Our coefficients of interest are the set of coefficients ω_s. For a skill category s, ω_s can be interpreted as the change in the wage-density elasticity associated with a one standard deviation increase in the relevant mean skill score K.

Table 9: Interactions Between Population Density and Mean Skill Scores

	(1) Log Real Wage
Population Density × Mean Communication Skill Score	-0.0154 (0.00836)
Population Density × Mean Managerial Skill Score	-0.00882 (0.00719)
Population Density × Mean Client-Facing Skill Score	0.0226*** (0.00663)
Population Density × Mean Quantitative Skill Score	-0.00747 (0.00819)
Population Density × Mean Analytical Skill Score	0.0481*** (0.00933)
Population Density × Mean Collaboration Skill Score	-0.0246*** (0.00680)
Population Density × Mean Physical Skill Score	-0.00286 (0.00546)
N	84317

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

As seen in Table 9, we estimate ω to be -0.015 for communication skills and -0.025 for collaboration skills. These statistically significant coefficients (albeit only at the 10% significance level for communication skills) suggest that occupations with more intensive use of the above skills enjoy a smaller urban wage premium. Conversely, we obtain positive and statistically significant estimates of ω for analytical and client-facing skills, at 0.0481 and 0.0226 respectively. Finally, we do not obtain statistically significant estimates of ω at any conventional significance levels for managerial, quantitative and physical skills. The biggest point of divergence between these results and those obtained from the main analysis is that more intensive use of communication skills is now associated with a lower urban wage premium. However, the results obtained here do lend further credence to our main results regarding how the urban wage premium is influenced by the intensity of use of analytical and client-facing skills.

6 Conclusion

Findings from our skill-based analysis of the urban wage premium support the notion that benefits arising from agglomeration economies are not uniformly distributed across occupations, but differ in substantial ways along the dimension of skills. A prominent pattern we find in our results is that occupations with skill profiles oriented towards more non-routine and job-specific skills enjoy greater benefits from agglomeration economies. This is consistent with our understanding of how agglomeration economies lead to productivity gains. It is, after all, occupations that demand a more complex set of skills which stand to gain more from search efficiencies and enhanced sharing and generation of knowledge and ideas.

We believe there is still much untapped potential in the skill-based analysis of agglomeration economies. One area yet to be comprehensively explored in the context of skill heterogeneity is how agglomeration economies facilitate learning and the accumulation of human capital. This points to the possibility of a skill-based analysis of the urban wage *growth* premium. Apart from this, the results we obtained from using more granular skill categories suggest there is more insight to be gained from applying alternative and perhaps more meticulous treatments of the categorisation of skills.

References

Andersson, M., J. Klaesson, and J. P. Larsson (2014): "The sources of the urban wage premium by worker skills: Spatial sorting or agglomeration economies?" *Papers in Regional Science*, 93, 727–747.

Bleakley, H. and J. Lin (2012): "Thick-market effects and churning in the labor market: Evidence from US cities," *Journal of urban economics*, 72, 87–103.

Cattell, R. B. (1966): "The scree test for the number of factors," *Multivariate behavioral research*, 1, 245–276.

Choi, T. (2020): "Agglomeration Effect of Skill-Based Local Labor Pooling: Evidence of South Korea," *Sustainability*, 12, 3198.

Combes, P.-P., G. Duranton, L. Gobillon, D. Puga, and S. Roux (2012): "The productivity advantages of large cities: Distinguishing agglomeration from firm selection," *Econometrica*, 80, 2543–2594.

D'Costa, S. and H. G. Overman (2014): "The urban wage growth premium: Sorting or learning?" *Regional Science and Urban Economics*, 48, 168–179.

De Blasio, G. and S. Di Addario (2005): "Do workers benefit from industrial agglomeration?" *Journal of regional Science*, 45, 797–827.

Di Addario, S. (2011): "Job search in thick markets," *Journal of Urban Economics*, 69, 303–318.

Duranton, G. and D. Puga (2004): "Micro-foundations of urban agglomeration economies," in *Handbook of regional and urban economics*, Elsevier, vol. 4, 2063–2117.

Elvery, J. A. (2010): "City size and skill intensity," *Regional Science and Urban Economics*, 40, 367–379.

Felstead, A., D. Gallie, F. Green, and G. Henseke (2019): "Skills and Employment Survey, 2017." [data collection]. UK Data Service. SN: 8581.

Gathmann, C. and U. Schönberg (2010): "How general is human capital? A task-based approach," *Journal of Labor Economics*, 28, 1–49.

Geel, R., J. Mure, and U. Backes-Gellner (2011): "Specificity of occupational training and occupational mobility: an empirical study based on Lazear's skill-weights approach," *Education Economics*, 19, 519–535.

Glaeser, E. L. and D. C. Mare (2001): "Cities and skills," *Journal of labor economics*, 19, 316–342.

Gould, E. D. (2007): "Cities, workers, and wages: A structural analysis of the urban wage premium," *The Review of Economic Studies*, 74, 477–506.

Graham, D. J. and P. C. Melo (2009): "Agglomeration economies and labour productivity: evidence from longitudinal worker data for GB's travel-to-work areas," .

Haller, P. and D. F. Heuermann (2020): "Opportunities and competition in thick labor markets: Evidence from plant closures," *Journal of Regional Science*, 60, 273–295.

Hirsch, B., E. J. Jahn, A. Manning, and M. Oberfichtner (2020): "The Urban Wage Premium in Imperfect Labor Markets," *Journal of Human Resources*, 0119–9960R1.

Horn, J. L. (1965): "A rationale and test for the number of factors in factor analysis," *Psychometrika*, 30, 179–185.

Kahn, J. H. (2006): "Factor analysis in counseling psychology research, training, and practice: Principles, advances, and applications," *The counseling psychologist*, 34, 684–718.

Kaiser, H. F. (1958): "The varimax criterion for analytic rotation in factor analysis," *Psychometrika*, 23, 187–200.

Lazear, E. P. (2009): "Firm-specific human capital: A skill-weights approach," *Journal of political economy*, 117, 914–940.

Marshall, A. (2009): *Principles of economics: unabridged eighth edition*, Cosimo, Inc.

Osterlind, S. J., B. G. Tabachnick, and L. S. Fidell (2001): *SPSS for Windows Workbook: To Accompany Tabachnick and Fidell" Using Multivariate Statistics"*, Allyn and Bacon.

Tabachnick, B. G., L. S. Fidell, and J. B. Ullman (2007): *Using multivariate statistics*, vol. 5, Pearson Boston, MA.

University of Essex, Institute for Social and Economic Research, NatCen Social Research, Kantar Public. (2020a): "Understanding Society: Waves 1-10, 2009-2019 and Harmonised BHPS: Waves 1-18, 1991-2009," [data collection]. 13th Edition. UK Data Service. SN: 6614.

——— (2020b): "Understanding Society: Waves 1-10, 2009-2019 and Harmonised BHPS: Waves 1-18, 1991-2009: Special Licence Access, Local Authority District," [data collection]. 12th Edition. UK Data Service. SN: 6666.

Yankow, J. J. (2006): "Why do cities pay more? An empirical examination of some competing theories of the urban wage premium," *Journal of Urban Economics*, 60, 139–161.

Zwick, W. R. and W. F. Velicer (1986): "Comparison of five rules for determining the number of components to retain." *Psychological bulletin*, 99, 432.

7 Appendix

7.1 Determination of Factorability of Data

In order to first determine the factorability of the data, we used the Kaiser-Meyer-Olkin (KMO) measure of sampling adequacy. The KMO measure "indicates the proportion of variance in the variables that might be caused by underlying factor" (IBM), and KMO values greater than 0.60 are required for good factor analysis (Osterlind et al., 2001) . We obtained a KMO value of 0.9317, which we treat as an affirmation that the SES2017 dataset is suitable for factor analysis.

7.2 Determination of Optimal Number of Factors

To determine the optimal number of factors (and thus skill categories), eigenvalues of each factor are calculated. Kaiser's criterion states that only factors with eigenvalues of at least one should be retained (Kahn, 2006; Kaiser, 1958), this corresponds to 10 factors as shown in Figure 1. However, Kaiser's criterion tends to overextract factors (Zwick and Velicer, 1986).

To explore the viability of using a smaller number of factors, we used Cattell's scree test. This involves examining a scree plot of eigenvalues for the factors and discerning the points at which significant drops occur (Cattell, 1966). From Figure 1, one can observe significant drops between factors one and two, factors four and five, and factors seven and eight. Cattel's scree test necessitates identifying the last significant drop in eigenvalues, thus the first seven factors are retained. A non-horizontal scree line can also be drawn through the point at which the eigenvalues flatten out (the solid green line in Figure 1), and factors above that scree line should be retained.

Horn's Parallel Analysis was also conducted to determine the optimal number of factors to retain (Horn, 1965). Parallel Analysis is a "sample-based adaptation of the population-based (Kaiser's) rule"(Zwick and Velicer, 1986) which uses random data simulation using the Monte Carlo Simulation Technique to determine the optimal number of factors.

From Figure 2, it can be seen that the adjusted Eigenvalue (that is, the difference between the unadjusted Eigenvalue, the actual data, and the Estimated Bias, the simulated data) becomes negative when shifting from the 9th factor to the 10th factor. As such, Parallel Analysis suggests that up to nine factors can be used for the factor analysis, since factors with unadjusted Eigenvalues greater than the random eigenvalues generated from the simulated data are able to explain more variance than chance (Kahn, 2006).

Given that these decision rules provided conflicting and convergent results, we further considered how our choices would affect the interpretability of factors and the simplicity of analysis. With these criteria in mind, we decided on seven factors - the smallest number from amongst our final options.

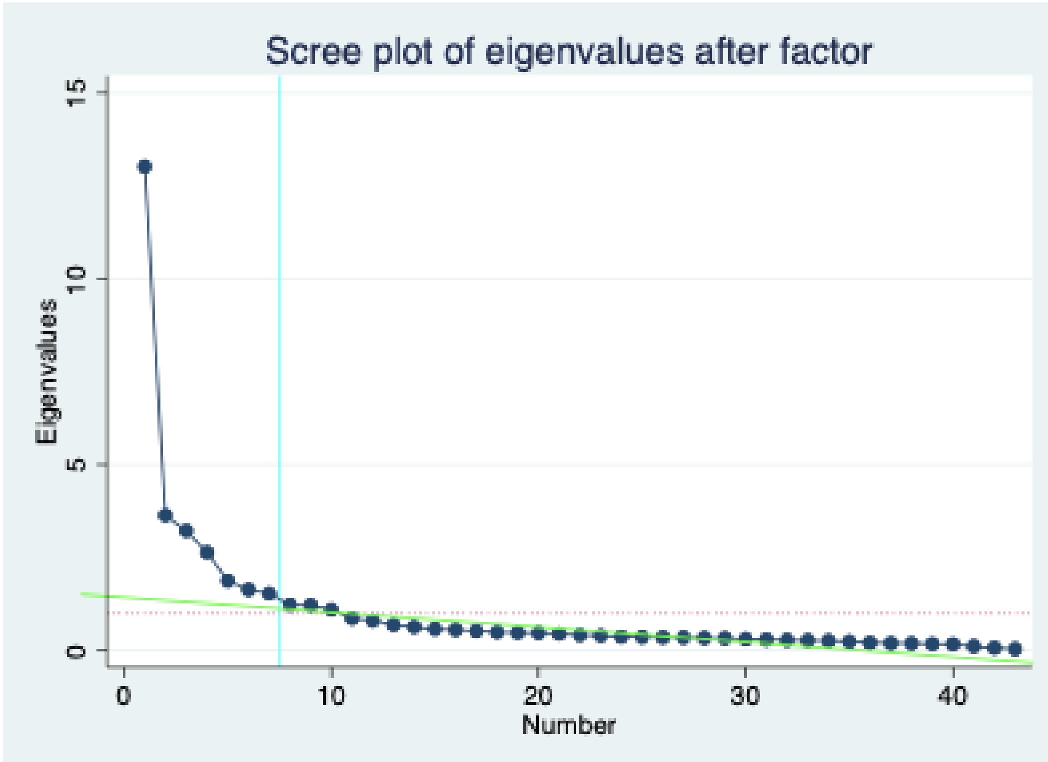

Figure 1: Scree Plot - Eigenvalue of the Last Factor Used vs Number of Factors Used
NOTE: The dashed red line indicates the point at which eigenvalues are higher than 1 (Kaiser's criterion). The solid green line indicates the line through the scree. The solid blue line indicates the point at which the eigenvalues drop significantly.

7.3 Interpretation of Factors

Factor 1: Communication Skills - communication skills are generally defined as the handling of information, specifically in receiving and delivering knowledge. Good communication skills result in clear, efficient and effective transmission of information and are thus vital and relevant to several occupations. As shown in Table 2, the skill categorization generated for the Communication Profile most notably includes reading and writing skills (all have factor loadings greater than 0.7). The ability to read written information and write materials is essential in transferring information, and as such, these skills are appropriately classed under the broader category of communication skills.

Factor 2: Managerial Skills - according to a study by Manktelow and Birkinshaw, surveying 15,242 managers worldwide, the most critical competencies of managers (in order of relevance) relate to maintaining working relationships, task management, decision making,

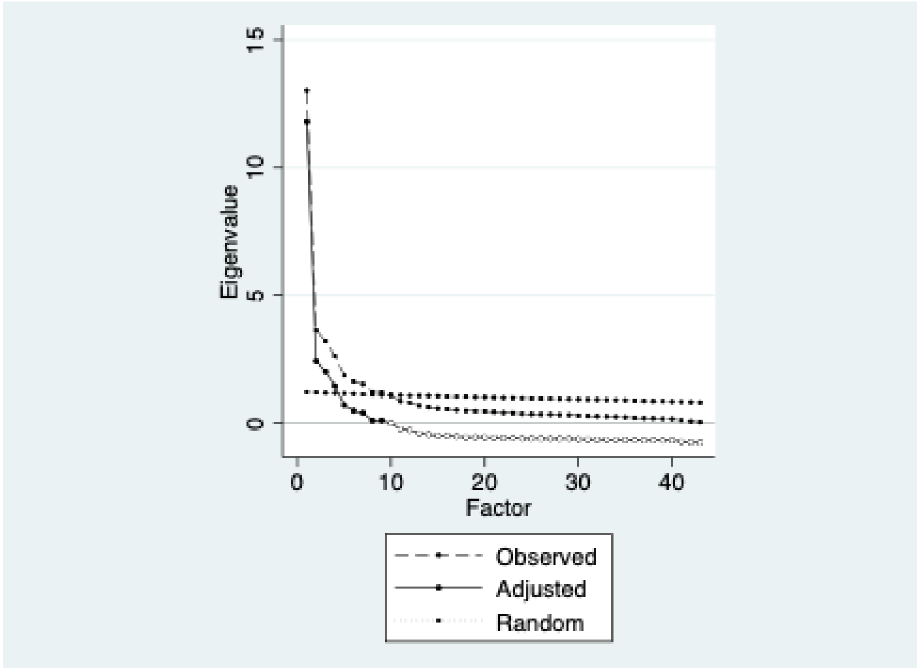

Figure 2: Eigenvalues of Actual data and simulated data

communication, understanding the needs of stakeholders, bringing people and resources together to problem solve, cultivating customer relations, building trust within the workforce and employing emotional intelligence. The specific skills with high factor loadings for the managerial profile correspond to "motivating staff", "keeping close control over resources", "coaching staff" "responsible for, developing careers of staff", "making strategic decisions about future" (all have factor loadings greater than 0.8); as the generally accepted skills required of managers encompass all of the skills corresponding to Factor 2, we interpret factor 2 as the category of managerial skills.

Factor 3: Analytical Skills - according to the Employer Skills Survey 2019 by the UK Department for Education, "complex analytical skills" are related to having specialist skills/knowledge, the ability to solve complex problems, knowledge of products/services and how the organisation works, the ability to read, write and understand instructions, and computer literacy and numerical skills. Seeing as how the skills with the highest factor loadings include "spotting problems or faults", "working out cause of problems/ faults", "thinking of solutions to problems" (all have factor loadings greater than or equal to 0.8)"analysing complex problems in depth" (factor loading exceeding 0.5), and correspond to these aforementioned "complex analytical skills", factor 3 is appropriately considered the "Analytical Skills" category.

Factor 4: Client-Facing Skills - The factor loadings for Factor 4 place the most emphasis on "looking the part", "sounding the part" (both have factors loading greater than 0.7) and "dealing with people" (factor loading greater than 0.6) and other similar skills that are closely linked to interaction with others, specifically clients. We therefore label Factor 4 the "Client-Facing Skills" category.

Factor 5: Quantitative Skills - according to Get Cover Letter, the top quantitative skills that professions require include mathematical, analytical, survey, and research skills. Considering how the skills with the highest factor loadings are "using a computer", "arithmetic" and "advanced mathematics" (all have factor loadings greater than 0.7), which are exceptionally relevant to the definition of quantitative skills, Factor 5 is suitably named the "Quantitative Skills" category.

Factor 6: Collaboration Skills - according to the University of Stathlyde Careers Services, collaboration skills relate to the giving and receiving of feedback, listening, communicating, supporting and acknowledging others' skills and contributions. Overall collaboration skills have to do with the "building and maintaining relationships" when "achieving the task". The factor loadings place the largest emphasis on "working with a team","listening carefully to colleagues","cooperating with colleagues" (all have factor loadings greater than 0.7) and thus Factor 6 is appropriately considered the "Collaboration Skills" category.

Factor 7: Physical Skills - according to Rhapsody Fitness, the top physical strength skills include endurance, stamina, strength, flexibility, speed, coordination, agility, balance, and accuracy. The skills with the highest factor loadings for Factor 7 include "Physical Strength", "physical stamina", "skill or accuracy in using hands/fingers", and "knowledge of use or operation of tools" (all have factor loadings greater than 0.7); accordingly, as the corresponding skill set fit well with the top physical skills defined above, Factor 7 is appropriately called the "Physical Skills" category.

Below we visualise the factor loadings of the most relevant skills for each skill category. The height of the bar represents the highest factor loading in the corresponding category. The top end of each colour portion represents the loading of the corresponding skill for that skill category.

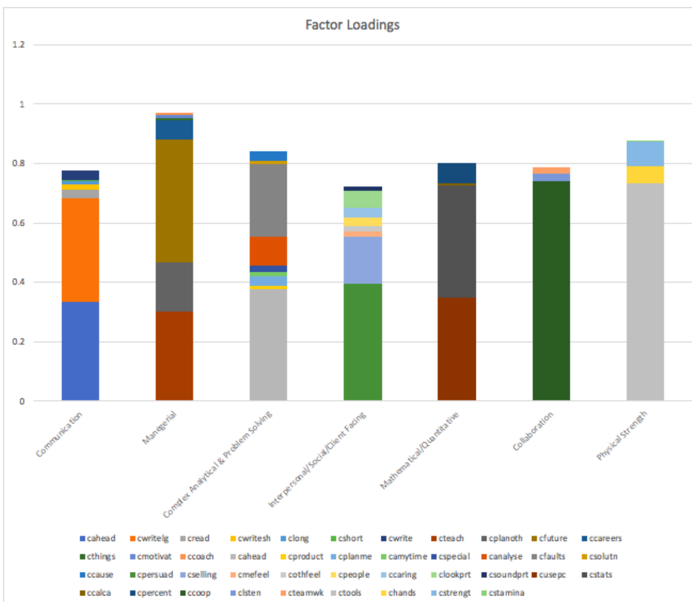

Figure 3: Visualisation of Factor Loadings by Skill Category

7.4 County-Level Filling Maps for Regional Skill Specialisations

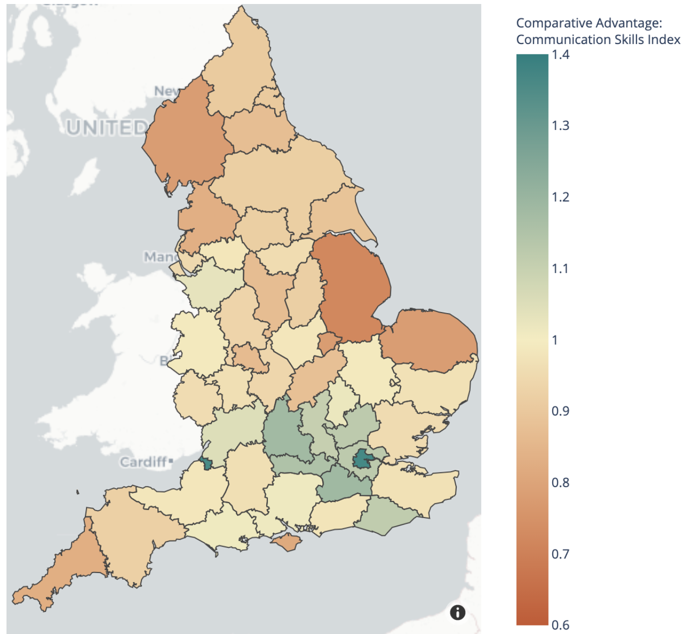

Figure 4: Filling Map for Communication Skills

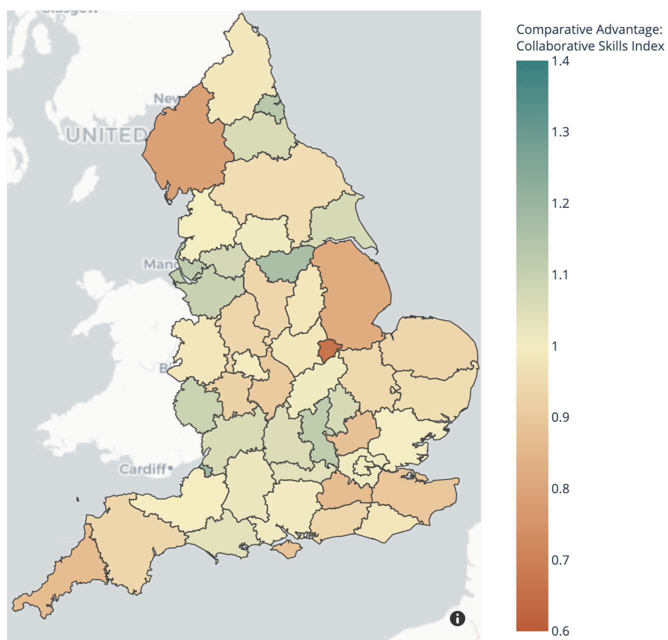

Figure 5: Filling Map for Collaboration Skills

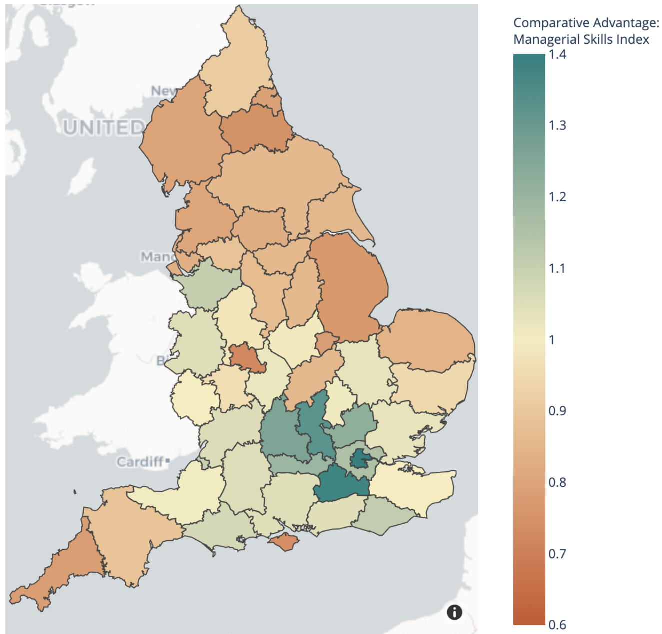

Figure 6: Filling Map for Managerial Skills

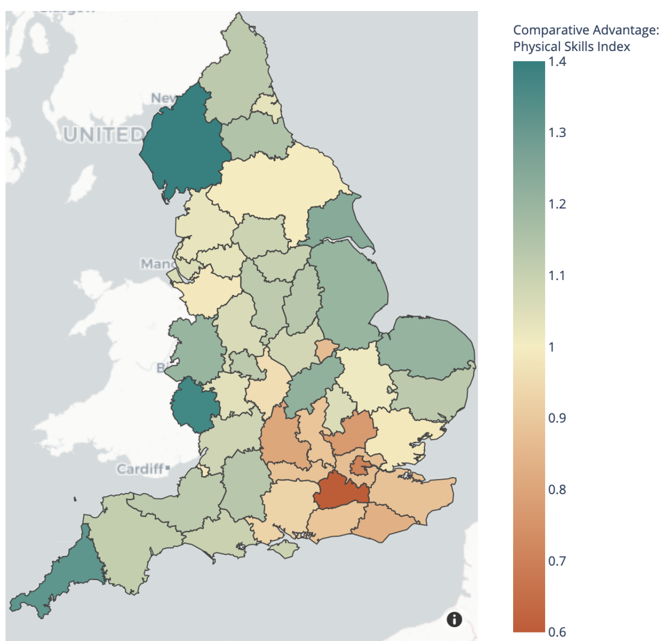

Figure 7: Filling Map for Physical Skills

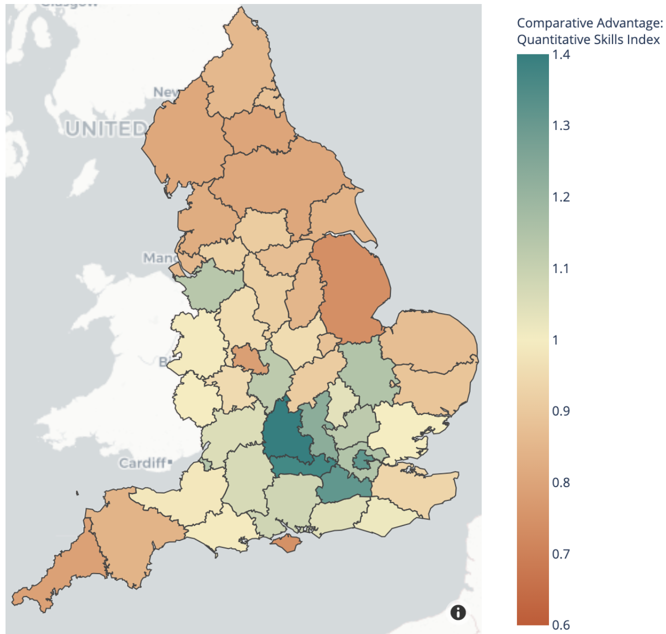

Figure 8: Filling Map for Quantitative Skills

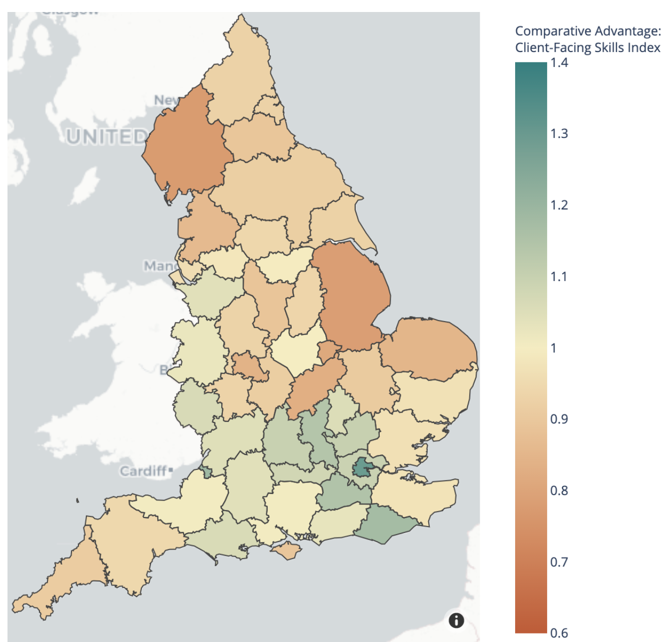

Figure 9: Filling Map for Client-Facing Skills

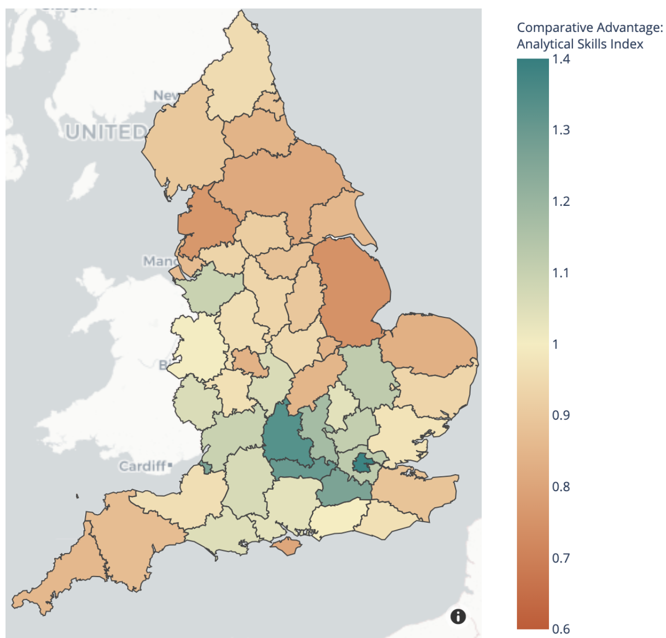

Figure 10: Filling Map for Analytical Skills