UCD CENTRE FOR ECONOMIC RESEARCH

WORKING PAPER SERIES

2022

Temporal Consistency of Forecasts and Data Releases

Constantin Bürgi, University College Dublin
Julio L Ortiz, Federal Reserve Board

WP22/07

February 2022

UCD SCHOOL OF ECONOMICS
UNIVERSITY COLLEGE DUBLIN
BELFIELD DUBLIN 4

Temporal Consistency of Forecasts And Data Releases*

Constantin Bürgi† Julio L. Ortiz‡

February 11, 2022

Abstract

We provide key insights on expectation formation based on the Bloomberg economic survey: around two thirds of professional forecasters provide GDP forecasts that are temporally consistent, meaning that quarterly forecasts add up to the annual. Temporally consistent forecasts are not more accurate than the inconsistent ones, but inconsistent ones might drive estimates of information frictions in some cases. For the overwhelming majority of consistent forecasts, annual GDP predictions almost immediately reflect the monthly GDP releases. These findings suggest that most forecasters make at least minor forecast updates after each data release. Indeed, the inattention rate is found to be between 3% and 6% at the quarterly frequency.

JEL: E37

Keywords: Temporal Forecast Reconciliation, Temporal Aggregation, Bloomberg Survey, GDP, Forecast Accuracy

---

*Declarations of interest: none

†*Email: constantin.burgi@ucd.ie*, University College Dublin, Belfield, Dublin 4, Ireland

‡*Email: julio.l.ortiz@frb.gov*, Federal Reserve Board, 20th and Constitution Avenue NW, Washington DC 20551. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

1 Introduction

Many economic decisions are dynamic in nature, making expectations formation critical for understanding economic phenomena such as the inflation, consumption-savings decisions, and firm investment. One approach to examine the drivers of expectations is to assess how economic relationships and accounting identities inform professional forecasts. So far, the literature on survey forecasts has mainly focused on relationships like Okun's Law (Mitchell and Pearce (2010), Pierdzioch et al. (2011) or Ball et al. (2015)), the Phillips Curve (Fendel et al. (2011)), and the Taylor rule (Mitchell and Pearce (2010)), concluding that forecasts tend to be more accurate if these relationships are taken into account.

Regarding accounting relationships, the literature has primarily studied the time series properties of forecasts across different frequencies or "temporal hierarchies", (see Rossana and Seater (1995) for a summary), and whether aggregate forecasts improve upon individual forecasts (e.g. Abraham (1982), Souza and Smith (2004), or Silvestrini et al. (2008)). More recently, Athanasopoulos et al. (2017) shows that consistent forecasting, defined as instances in which high frequency predictions aggregate up to low frequency predictions, can improve forecasts when there is model uncertainty. However, to our knowledge, examining whether participants in forecasting surveys take into account the temporal hierarchy, and whether that improves the predictions, remains an open question. In this paper, we compare quarterly and annual predictions made by the same forecaster in the same period, assess the frequency with which quarterly predictions "aggregate up" to the annual forecast, and document facts relating to consistent and inconsistent forecasters.

So far, checking whether survey forecasts are temporally consistent has been a challenge since many surveys do not provide sufficient information to test this. For instance, the Survey of Professional Forecasters from the European Central Bank, Consensus Economics or BlueChip only provide one frequency (annual), or the Survey of Professional Forecasters from the Philadelphia Fed (SPF) imposes temporal consistency. In this paper, we take advantage of the US GDP predictions in the Bloomberg economic survey (BBG) which asks for predictions at both the quarterly and annual frequency and does not impose temporal consistency.[1]

In addition to testing for temporal consistency, we find that the temporal hierarchy provides a key insight into the updating behavior of forecasters. Right after a quarterly GDP release, forecasts only remain temporally consistent if they either do not bring in the new release or if they immediately update their prediction as well. Comparing the forecasts that are consistent with the new release to the ones that are consistent with the data prior to the release provides a new way to estimate the the attention rate of forecasters and hence the extent of sticky information (see Mankiw and Reis (2002), Clements (2012), Andrade and Le Bihan (2013) or Coibion and Gorodnichenko (2015)).[2]

Because forecasters that are typically consistent might have a different expectations formation process from forecasters that predict quarterly and annual GDP independently, we also compare the forecast efficiency (e.g. see Nordhaus (1987)) and whether there are any information frictions (e.g. see Coibion and Gorodnichenko (2015)) for both types of forecasters separately. For example, if most temporally consistent forecasts update immediately after data releases, this would imply that sticky information (Mankiw and Reis, 2002) cannot be the driver for a potential underadjustment to new information for this type of forecaster.

The remainder of the paper is structured as follows: The next section describes the data, followed by the empirical section and forecast accuracy. We next look at forecast efficiency and how temporal consistency changes over time. The last section concludes.

2 Data

We use US GDP predictions from the BBG survey from 2000Q2 to 2015Q1. We focus on forecasts made in the last month of each quarter (March, June, September and December).

[1]The Wall Street Journal economic survey (WSJ) offers another way to test this for GDP but its annual prediction is defined as the fourth quarter over fourth quarter change. Estimates of temporally consistent forecasts for the WSJ survey are provided in the appendix and are in line with the Bloomberg results.

[2]This is also closely related to the literature on revisions around large shocks like Baker et al. (2020) and Gaglianone et al. (2020).

This survey asks both the annualized quarterly path and the annual GDP number both in terms of growth rates. The survey is released at the end of a month, which means that the third release of GDP is known at the time of the survey. Unlike the SPF, the BBG survey does not impose quarterly-to-annual forecast consistency. The participants of the survey are professional forecasters who often belong to economics departments of banks, firms or research institutions and publish these forecasts regularly.

Since a part of our analysis requires constructing prediction errors, we also make use of GDP releases. For the quarterly predictions, we take the real-time (first) GDP release. For the annual prediction, we use the January GDP release.

3 Temporally Consistent Forecasts And Updating

In order to test the temporal consistency of predictions, we opt to use the quarterly predictions and infer an implied annual GDP growth rate.[3] For the appropriate calculation of the implied = annual $CY\%\Delta$, eight quarterly levels are necessary (or seven growth rates). Up to four of the quarterly levels are predictions that need to be converted into levels and the remaining ones are based on the real-time data releases. The annual year over year number for the current year is then calculated by first averaging the quarterly levels for the current and previous year and then calculating the percentage change between the two averages. Specifically,

$$CY\%\Delta = \frac{\sum_{i=1}^{4} Q_i L}{\sum_{i=1}^{4} Q_i LL} * 100 - 100 \tag{1}$$

where $Q_i L$ is the predicted/most recent level for the current year quarter $i \in \{1, 2, 3, 4\}$ and $Q_i LL$ is the realized most recent level for quarter $i \in \{1, 2, 3, 4\}$ last year. The levels are in turn constructed based on the real-time levels and the levels implied by the predicted

[3]Theoretically, one could use a combination of the annual and all but one of the quarterly predictions to calculate the implied prediction for the remaining quarter. However, as the quarterly values are reported in terms of annualized growth, they have more precision than the annual prediction. In turn, the implied quarterly prediction would be less precise and as a result likely identify more consistent predictions than our preferred method.

quarterly growth rates. The implied annual prediction is then rounded to the next tenth of a percent to match the precision of the annual predictions.

Because temporally consistent forecasts might use different data releases, this leads to a large number of potential combinations. For example, one prediction at the end of March might be based on the third release of Q4 GDP while another might be based on the second release of Q4 GDP. The survey is released shortly after the third release of GDP and as a result, most temporally consistent forecasts will be based on either the second or third release of GDP. We thus deem an annual forecast temporally consistent, if its quarterly predictions together with either the second or the third release of GDP imply the same annual prediction rounded to the same precision.

Table 1 shows the fraction of forecasters that are temporally consistent based on either the second or third release for each quarter. The first row shows that around two thirds of forecasters produce temporally consistent forecasts. The third row in the table shows that very few forecasts are in line with the (outdated) second release, and together with the second row it becomes clear that they are also in line with the third release most of the time.[4] This suggests that almost all temporally consistent forecasters are also fully attentive and update their predictions immediately after a data release. Indeed, the lower bound for the inattention based on the consistent forecasts would be 2.8% and the upper bound 6.3%, in the range of Andrade and Le Bihan (2013). These are calculated by taking the fraction of forecasters consistent with the second release relative to the ones consistent with either the third and second release. The range arises because the cases where both types of consistent forecasts are identical can be attributed to either side.

If these numbers are taken at face value, they have some important implications for the expectation formation of agents. Specifically, different forecasters might have different updating behaviors. A temporally consistent forecaster must always update the quarterly numbers together with the annual numbers. This is not necessarily the case for forecasters

[4] Given the low number of forecasts consistent with the second most likely scenario (the second release), other scenarios should have even lower numbers.

Table 1: Share of Temporally Consistent Forecasts

Month	3	6	9	12
2nd or 3rd Release	0.626	0.644	0.683	0.738
3rd Release	0.600	0.632	0.668	0.719
2nd Release	0.069	0.017	0.039	0.038
N	962	881	753	524

Note: This table takes the latest available data for each forecaster and shows the fraction of forecasts for which the annual year over year GDP growth (forecast) is identical to the annual GDP growth implied by taking the quarterly predictions. For months 1-4, all 4 quarters are forecasts. For months 5-7 Q2-Q4 are forecasts and Q1 uses the latest actual. For months 8-10, Q1 and Q2 are actual data and for the remaining months, Q1-Q3 are the actual data.

that predict the quarterly and annual numbers independently from each other. As temporally consistent forecasts also tend to immediately update after data releases, there is also an implication for information frictions. The only way how substantial stickiness or inattention can enter expectation formation is through forecasts that are not temporally consistent. Note that there could still be noisy information, forecast smoothing or other factors causing a positive weight on past predictions.

A key assumption here is that forecasters actively choose whether to be temporally consistent or not. An alternative explanation could be that the temporal inconsistencies are due to errors or mistakes in the survey or some other constraints (for example, the SPF requires forecasters to be temporally consistent) which we address next.

4 Are Consistent Forecasts Better?

So far, forecasters were assumed to choose whether to independently forecast quarterly and annual predictions or satisfy temporal consistency. An immediate implication of this assumption is that imposing consistency on all forecasts should worsen the prediction for the forecasters that independently forecast the two frequencies. Indeed, the period average of inconsistent annual forecasts has a mean squared prediction error of 0.119 and the matched sample of the period average of the implied annual forecasts based on the quarterly predictions is 0.151 or 27% higher. While this difference is substantial and in line with the assumption, it is not statistically significant based on the two-sided Diebold and Mariano (1995) test with the adjustments in Harvey et al. (1997) for squared errors (p-value of 0.27). One implication of this result is that surveys that require forecasts to be temporally consistent (e.g. the SPF) make some forecasters perform worse than they otherwise would. Based on the Bloomberg survey, this could affect a third of forecasts and thus the overall average as well. Indeed, the overall average is 8% worse when the consistency constraint is imposed (also not statistically significant).

Another implication of the choice assumption and hence forecast efficiency is that temporally consistent forecasters should be no better than forecasters that forecast the frequencies independently. While temporally consistent forecasting can lead to more accurate forecasts (e.g. Athanasopoulos et al. (2017)), this depends on the data generating process and is not always the case. Indeed, if one of the two methods of forecasting was clearly better, rational forecasters would all choose that method. In order for some rational forecaster to choose to predict the frequencies independently and others to choose temporally consistent forecasting methods, neither method should be superior to the other.

In order to test whether one method produces more accurate predictions than the other, we first need to define who is a temporally consistent forecaster and who is not. Based on the share of consistent forecast in the first quarter across all years, we determine whether a forecaster is typically consistent or not. For each period, we calculate the average forecast for the annual prediction (YoY), the current quarter prediction (CQ), the one quarter ahead prediction (1Q), the two quarter ahead forecast (2Q) and the three quarter ahead forecast (3Q) for the two groups separately and use the first release to determine the prediction errors. Table 2 shows the relative performance of the two groups. We used three thresholds to determine typically consistent forecast (70%, 80% and 90%) and report the relative mean squared prediction errors (MSE) of the groups. That is, we divide the MSE of the typically consistent group by the MSE of the inconsistent group. We also run a two-sided Diebold and Mariano (1995) test with the adjustments in Harvey et al. (1997) for squared errors to test for significance.

Table 2: Mean Squared Errors Across Thresholds and Horizons BBG

Threshold	YoY	CQ	1Q	2Q	3Q
0.7	1.034	0.986	1.013	1.010	0.962
0.8	1.036	0.976	1.032	1.007	0.958
0.9	1.193**	0.964	1.021	1.042	1.000

Note: The table shows the MSE of the consistent forecasts relative to the inconsistent ones. The stars are based on a two-sided Diebold-Mariano test with squared loss function. *, **, *** imply significant difference at 10%, 5% and 1% level.

In line with the assumption that one method should not be superior to the other, there is no clear pattern in regards to the MSE. Indeed, in some cases the temporally consistent forecasts are more accurate and sometimes the inconsistent ones are and only one case is statistically significant.

5 Forecast Efficiency

Next, we run a Nordhaus (1987) and Coibion and Gorodnichenko (2015) type test to assess whether temporal consistency has any impact on information friction and forecast efficiency. Specifically, we run the regression

$$FE_{t,t+h} = \alpha + \beta FR_{t,t+h} + \varepsilon_t \tag{2}$$

where $FE_{t,t+h}$ is the prediction error corresponding to forecasts made in period t for period $t+h$ minus the first release of the actual for period $t+h$. $FR_{t,t+h}$ is the forecast revision corresponding to the prediction made in period t for period $t+h$ minus the prediction made in period $t-1$ for period $t+h$. If forecasts are efficient, one would expect no relation between prediction errors and revisions. If there are information frictions or forecasters smooth their predictions, one would expect a positive coefficient. A negative coefficient would imply that forecasters overadjust their predictions.

Table 3: Current Quarter - 2 Quarter Ahead Forecast Efficiency

	Dependent variable:						
	Prediction Error						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Threshold	0.7		0.8		0.9		
Group	Cons	Indep	Cons	Indep	Cons	Indep	All
CQrev	0.132	0.287**	0.112	0.263**	0.068	0.243**	0.206**
	(0.087)	(0.115)	(0.078)	(0.119)	(0.076)	(0.110)	(0.101)
Q1rev	−0.059	0.073	−0.124	0.095	−0.057	0.023	0.018
	(0.364)	(0.437)	(0.357)	(0.431)	(0.327)	(0.414)	(0.403)
Q2rev	−0.193	0.308	−0.484***	0.269	−0.254	0.157	0.087
	(0.798)	(0.512)	(0.027)	(0.557)	(0.670)	(0.699)	(0.702)
Obs	59	59	59	59	59	59	59

Note: The table shows the coefficient of interest for 21 regressions and each row uses a different horizon. The Cons columns use the period averages for consistent forecasters based on the three thresholds and the Indep columns use the period averages for inconsistent forecasters for the same thresholds. The column All includes all forecasters in the average. *, **, *** imply significant difference at 10%, 5% and 1% level based on Newey-West standard errors.

Table 3 shows the coefficient of interest for various cuts of the data. The first two columns show the coefficient for the period average of forecasters that are at least 70% of the time temporally consistent (Cons) or not (Indep) and three horizons. The next two groups of columns repeat the regressions with thresholds of 80% and 90%, respectively. The last column is the regression for the period averages of all forecasters.

The current quarter predictions show that temporally consistent forecasters are efficient and the temporally inconsistent forecasters underadjust their forecasts. This underadjustment can be driven by forecast smoothing or some information frictions. The last column shows that the underadjustment of temporally inconsistent forecasters is driving the underadjustment at the aggregate level. For one quarter and two quarter ahead, the picture changes a bit with consistent forecasts having a negative coefficient but mostly insignificant, implying an overadjustment. Inconsistent forecasters and the overall aggregate show no deviations from efficient forecasting. The overadjustment for consistent forecasters is of particular interest due to individual forecaster level overadjustment documented for example in Bordalo et al. (2020) or Kohlhas and Walther (2021).

6 Temporal Consistency Over Time

Another interesting aspect about temporal consistency is whether it changed over time. Instead of splitting the sample by quarter in Table 1, we split the sample by year in this section.

As Figure 1 shows, there is no general upward or downward trend in the share of temporally consistent predictions over the years. It is the case that a low in the share of consistent forecasts was reached in 2003 and there was a subsequent increase until 2010 but this was not sustained. After 2010, the share of temporally consistent forecasters declines again. The observed pattern could be part of a cycle, particularly due to the steady increase and decline in the share of temporally consistent predictions. While peaks in the share were reached shortly after recessions 2002 and 2010 there is not enough data ensure that this pattern is related to the business cycle.

In contrast to the share of temporally consistent forecasts, the number of consistent

Figure 1: Share of Forecasts Consistent With The Latest Release

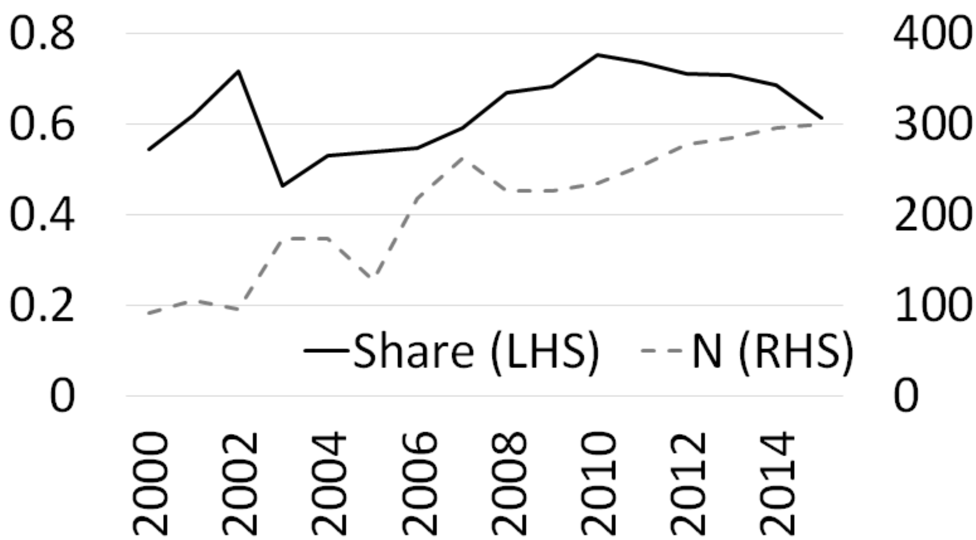

forecasts is steadily increasing, as is the number of forecasters in the survey overall.[5]

7 Conclusion

In this paper, we show that about two thirds of the participants in the Bloomberg survey provide temporally consistent forecasts and about one third updates quarterly and annual predictions independently from each other. We further show that the overwhelming majority of consistent forecasts immediately reflect new data releases. This suggests that inattention is minimal. Temporally inconsistent forecasts were found to be no worse than consistent forecasts and imposing consistency adversely affects forecast accuracy. These two properties are in line with the assumption that the temporally inconsistent forecasts do not merely reflect errors and mistakes. This suggests that imposing temporal consistency might have adverse effects on aggregate forecast.

We also compare forecast efficiency across types of forecasters. It was found that temporally consistent forecasters tend to place a lower weight on past predictions. In some cases, this results in aggregate underadjustments being purely driven by inconsistent forecasters.

[5] The 2015 number has been annualized as our data only includes the first quarter.

In turn, future research might be able to assess whether the overadjustment of the average consistent forecasters is related to the overadjustment of individual forecasters or not.

References

Abraham, B. (1982). Temporal aggregation and time series. *International Statistical Review/Revue Internationale de Statistique*, pages 285–291.

Andrade, P. and Le Bihan, H. (2013). Inattentive professional forecasters. *Journal of Monetary Economics*, 60(8):967–982.

Athanasopoulos, G., Hyndman, R. J., Kourentzes, N., and Petropoulos, F. (2017). Forecasting with temporal hierarchies. *European Journal of Operational Research*, 262(1):60–74.

Baker, S. R., McElroy, T. S., and Sheng, X. S. (2020). Expectation formation following large, unexpected shocks. *Review of Economics and Statistics*, 102(2):287–303.

Ball, L., Jalles, J. T., and Loungani, P. (2015). Do forecasters believe in okun's law? an assessment of unemployment and output forecasts. *International Journal of Forecasting*, 31(1):176–184.

Bordalo, P., Gennaioli, N., Ma, Y., and Shleifer, A. (2020). Overreaction in Macroeconomic Expectations. *American Economic Review*, 110:2748–2782.

Clements, M. P. (2012). Do professional forecasters pay attention to data releases? *International Journal of Forecasting*, 28(2):297–308.

Coibion, O. and Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–78.

Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics*, 13(3):253–263.

Fendel, R., Lis, E. M., and Rülke, J.-C. (2011). Do professional forecasters believe in the phillips curve? evidence from the g7 countries. *Journal of Forecasting*, 30(2):268–287.

Gaglianone, W. P., Giacomini, R., Issler, J. V., and Skreta, V. (2020). Incentive-driven inattention. *Journal of Econometrics*.

Harvey, D., Leybourne, S., and Newbold, P. (1997). Testing the equality of prediction mean squared errors. *International Journal of forecasting*, 13(2):281–291.

Kohlhas, A. N. and Walther, A. (2021). Asymmetric Attention. *American Economic Review*, 111:2879–2925.

Mankiw, N. G. and Reis, R. (2002). Sticky information versus sticky prices: a proposal to replace the new keynesian phillips curve. *The Quarterly Journal of Economics*, 117(4):1295–1328.

Mitchell, K. and Pearce, D. K. (2010). Do wall street economists believe in okun's law and the taylor rule? *Journal of Economics and Finance*, 34(2):196–217.

Nordhaus, W. D. (1987). Forecasting efficiency: concepts and applications. *The Review of Economics and Statistics*, pages 667–674.

Pierdzioch, C., Rülke, J.-C., and Stadtmann, G. (2011). Do professional economists' forecasts reflect okun's law? some evidence for the g7 countries. *Applied Economics*, 43(11):1365–1373.

Rossana, R. J. and Seater, J. J. (1995). Temporal aggregation and economic time series. *Journal of Business & Economic Statistics*, 13(4):441–451.

Silvestrini, A., Salto, M., Moulin, L., and Veredas, D. (2008). Monitoring and forecasting annual public deficit every month: the case of france. *Empirical Economics*, 34(3):493–524.

Souza, L. R. and Smith, J. (2004). Effects of temporal aggregation on estimates and forecasts of fractionally integrated processes: A monte-carlo study. *International Journal of Forecasting*, 20(3):487–502.

A Temporally Consistent Forecasts WSJ

The main text focused on the analysis for the Bloomberg survey. In this appendix, we repeat the consistency analysis for the monthly Wall Street Journal Survey. Similarly to the BBG survey, we use the quarterly path to impute an annual prediction which we then compare to the actual annual prediction. The WSJ survey uses the fourth quarter over fourth quarter growth rate for the annual. Thus the imputed annual becomes

$$CY\%\Delta = \prod_{i=1}^{4}(1 + Q_iL/100)^{1/4} * 100 - 100 \tag{3}$$

where Q_iL is the quarterly growth prediction/release for quarter i. As with the Bloomberg survey, the imputed value uses either the prediction or the latest available growth rate. Table 4 shows the share of consistent forecasters in each month.

Table 4: Temporally Consistent Q4/Q4 Forecasts

Month	1	2	3	4	5	6	7	8	9	10	11	12
Share	0.604	0.643	0.614	0.593	0.563	0.610	0.599	0.570	0.586	0.659	0.577	0.619
N	646	680	666	619	670	724	654	688	705	636	683	709

Note: The table takes the latest available data for each forecaster and shows the fraction of forecasts for which the Q4/Q4 GDP growth forecast is within 0.01 of the Q4/Q4 GDP growth implied by taking the quarterly forecast. For months 1-4, all 4 quarters are forecasts. For months 5-7 Q2-Q4 are forecasts and Q1 uses the latest actual. For months 8-10, Q1 and Q2 are actual data and for the remaining months, Q1-Q3 are the actual data.

As Table 4 shows, the consistency measures are around 60%, slightly lower than the Bloomberg survey.

UCD CENTRE FOR ECONOMIC RESEARCH – RECENT WORKING PAPERS

WP21/10 Benjamin Elsner, Ingo E Isphording and Ulf Zölitz: 'Achievement Rank Affects Performance and Major Choices in College' April 2021
WP21/11 Vincent Hogan and Patrick Massey: 'Soccer Clubs and Diminishing Returns: The Case of Paris Saint-Germain' April 2021
WP21/12 Demid Getik, Marco Islam and Margaret Samahita: 'The Inelastic Demand for Affirmative Action' May 2021
WP21/13 Emmanuel P de Albuquerque: 'The Creation and Diffusion of Knowledge - an Agent Based Modelling Approach' May 2021
WP21/14 Tyler Anbinder, Dylan Connor, Cormac Ó Gráda and Simone Wegge: 'The Problem of False Positives in Automated Census Linking: Evidence from Nineteenth-Century New York's Irish Immigrants' June 2021
WP21/15 Morgan Kelly: 'Devotion or Deprivation: Did Catholicism Retard French Development?' June 2021
WP21/16 Bénédicte Apouey and David Madden: 'Health Poverty' July 2021
WP21/17 David Madden: 'The Dynamics of Multidimensional Poverty in a Cohort of Irish Children' August 2021
WP21/18 Vessela Daskalova and Nicolaas J Vriend: 'Learning Frames' August 2021
WP21/19 Sanghamitra Chattopadhyay Mukherjee: 'A Framework to Measure Regional Disparities in Battery Electric Vehicle Diffusion in Ireland' August 2021
WP21/20 Karl Whelan: 'Central Banks and Inflation: Where Do We Stand and How Did We Get Here?' August 2021
WP21/21 Tyler Anbinder, Cormac Ó Gráda and Simone Wegge: '"The Best Country in the World": The Surprising Social Mobility of New York's Irish Famine Immigrants' August 2021
WP21/22 Jane Dooley and David Madden: 'Ireland's Post Crisis Recovery, 2012-2019: Was It Pro-Poor?' September 2021
WP21/23 Matthew Shannon: 'The Impact of Victimisation on Subjective Well-Being' September 2021
WP21/24 Morgan Kelly: 'Persistence, Randomization, and Spatial Noise' October 2021 *(For revised version of this paper see WP21/25)*
WP21/25 Morgan Kelly: 'Persistence, Randomization, and Spatial Noise' November 2021
WP21/26 Eliane Badaoui and Frank Walsh: 'Productivity, Non-Compliance and the Minimum Wage' November 2021
WP21/27 Annette Broocks and Zuzanna Studnicka: 'Gravity and Trade in Video on Demand Services' December 2021
WP21/28 Linda Mastrandrea: 'Linking Retail Electricity Pricing and the Decarbonisation of the Energy Sector: a Microeconomic Approach' December 2021
WP22/01 Doina Caragea, Theodor Cojoianu, Mihai Dobri, Andreas Hoepner and Oana Peia 'Competition and Innovation in the Financial Sector: Evidence from the Rise of FinTech Start-ups' January 2022
WP22/02 Sandra E Black, Paul J Devereux, Fanny Landaud and Kjell G Salvanes: 'The (Un)Importance of Inheritance' January 2022
WP22/03 Claes Ek and Margaret Samahita: 'Pessimism and Overcommitment: An Online Experiment with Tempting YouTube Content' January 2022
WP22/04 Paul J Devereux: 'Fragility of the Marginal Treatment Effect' January 2022
WP22/05 Lucie Martin, Liam Delaney and Orla Doyle: 'Everyday Administrative Burdens and Inequality' February 2022
WP22/06 Karl Whelan: 'The Past, Present and Future of Euro Area Monetary-Fiscal Interactions' February 2022

UCD Centre for Economic Research
Email economics@ucd.ie