PARETO

Undergraduate Journal of New Economists

University of Toronto Mississauga

2023 Issue

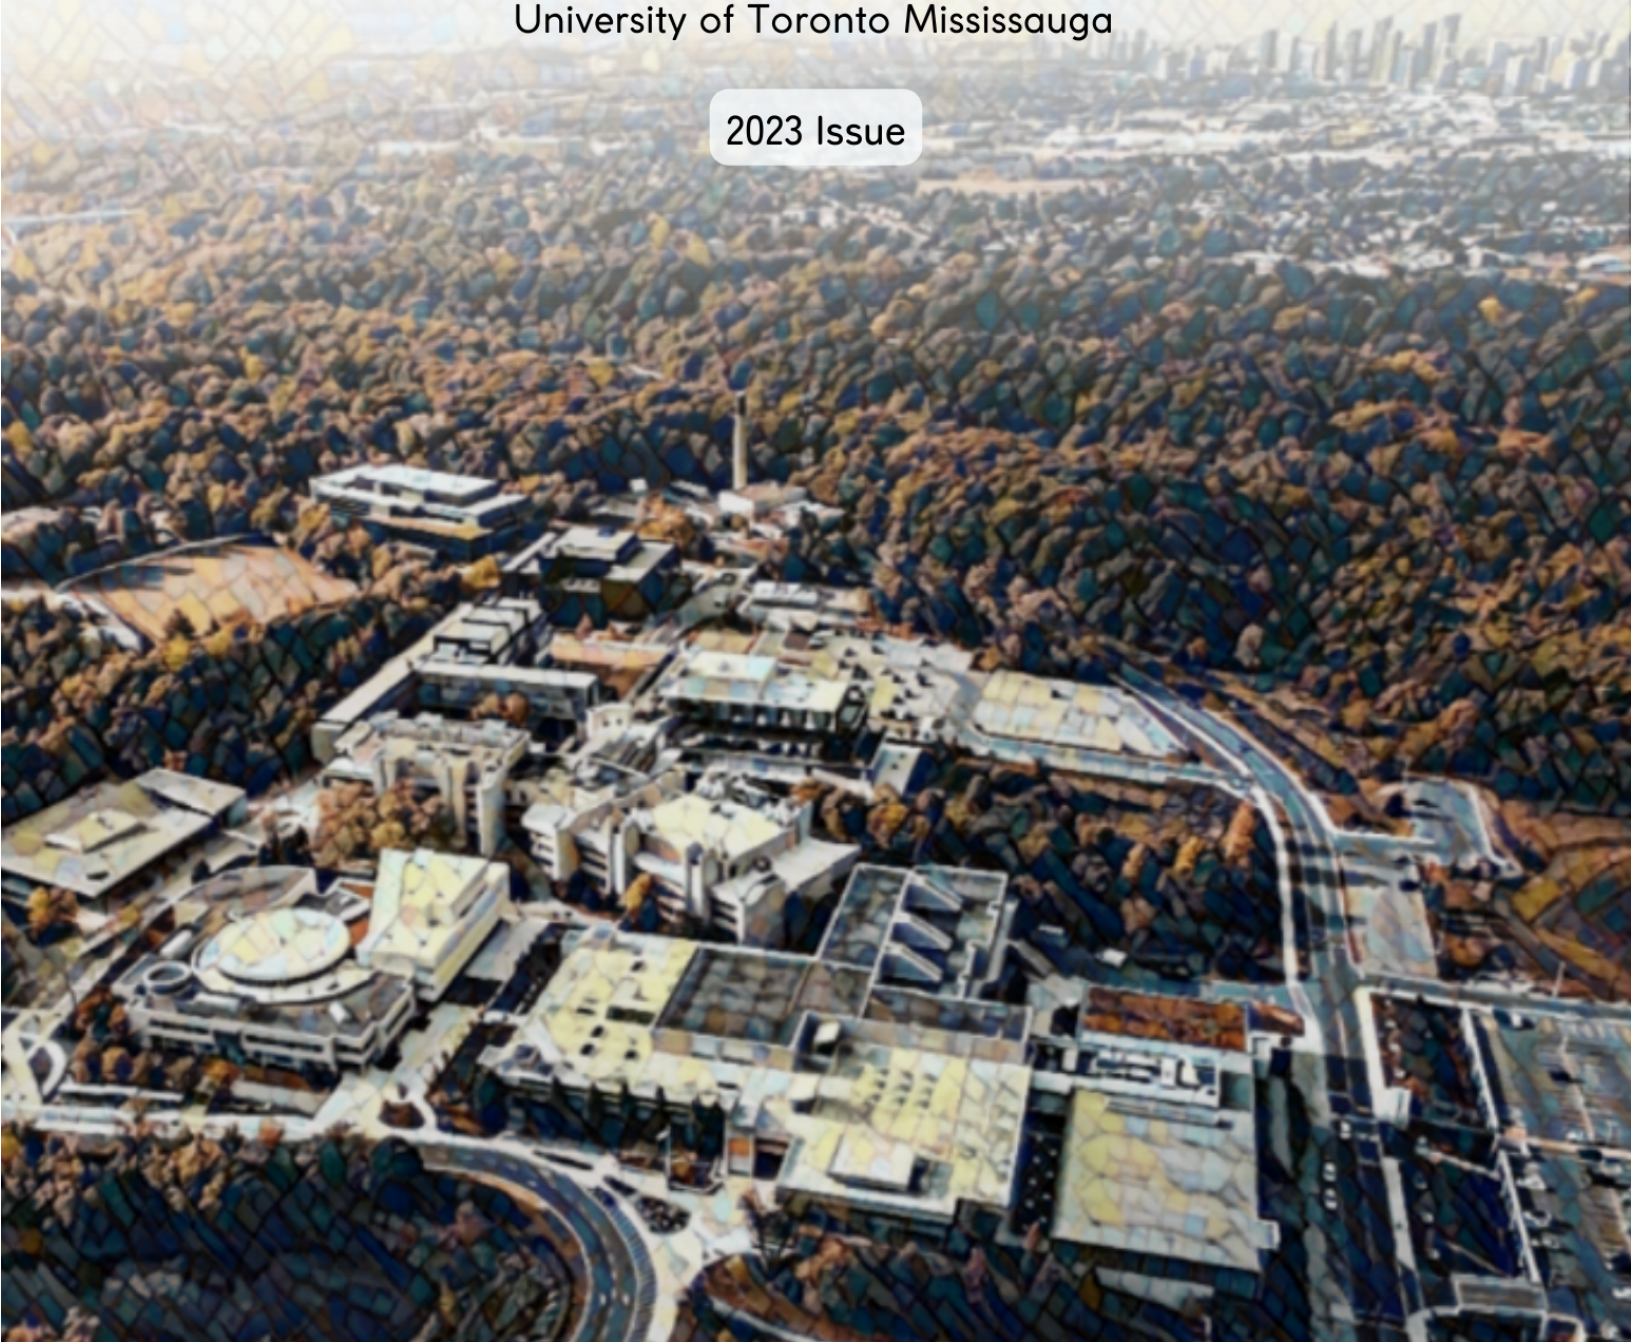

Table of Contents

1 Front Matter *by Nicholas Zammit*

Forward, pp. 1.

2 Effects of Cannabis Legalization on Tobacco and Alcohol Prevalence *by Gregory Chung*

Introduction, pp. 3 • The Extant Literature, pp. 4 • Data & Methods, pp. 5 • Results, pp. 8 • Conclusions, pp. 13.

3 Skilled Immigration and its Impact on Ontarios Unemployment Rates *by Ravdeep Sandal*

Introduction, pp. 19 • Literature Review, pp. 20 • Summary of the Canadian Immigration System, pp. 21 • Data, pp. 23 • Model, pp. 25 • Results, pp. 26 • Conclusion, pp. 30.

4 Estimating the "Breadwinner Impact" on Household Division of Labour *by Melissa Siqueira*

Introduction, pp. 37 • Literature Review, pp. 37 • Data and Methodology, pp. 39 • Results, pp. 44 • Conclusions, pp. 49 • Implications, pp. 51.

5 The Effects of Unconventional Monetary Policy on Firm Capital Structure *by Jun Takahashi*

Introduction, pp. 56 • Existing Literature, pp. 57 • Data Summary, pp. 59 • Model Design, pp. 63 • Results, pp. 65.

6 Modeling Social Learning Using Dyna-Q and Ant Colony Optimization *by Renato Zimmermann*

Introduction, pp. 86 • Related Work & Motivation, pp. 87 • Model, pp. 91 • Simulations, pp. 94 • Conclusions, pp. 98.

PARETO

Undergraduate Journal of New Economists

University of Toronto Mississauga

Forward

It is with great enthusiasm and pride that I am able to introduce the first edition of Pareto. Pareto-efficiency is a term familiar to many students of economics. When resources are allocated in the most economically efficient way they are considered to be in a Pareto-efficient state. Another familiar term will be Pareto-optimality. Pareto-optimality holds a normative tone describing a situation where no further reallocations can make anyone better off without making anyone else worse off. For our journal the term Pareto contextualizes both of these definitions. Students are encouraged to make the university community, the field of study and broader world better off through their research. The contributions made by the students in these articles can only add to our collective understanding and encourage greater curiosity. In the very least their pursuit of knowledge and deep exploration into a subject has made them better researchers and stronger contributors to the academic community.

From my perspective, Pareto represents the fulfillment of a promise made to showcase top undergraduate work that was generated in my ECO420 Research in Applied Economics course. Having taught undergraduate dissertation courses at both the University of Cambridge and the University of Warwick, I recognized in our students a similar desire to engage with the research process. In the first year of the course we had six students produce very high quality undergraduate dissertations. Amongst them were two of the papers selected for the first edition of Pareto. Successive cohorts have continued to impress and a myriad of quality work has stockpiled. Most recently one of the research articles produced for ECO420 was able to gain acceptance to the prestigious Carroll Round undergraduate dissertation competition at Georgetown University. In such a short time our students have been able to place the University of Toronto Mississauga on the map in terms of undergraduate research in economics.

Pareto currently stands as the only undergraduate economics journal across the three campuses at the University of Toronto. Judging by the demand for Research Opportunity Projects and

fourth year reading courses, I would say our students have an insatiable hunger to participate in undergraduate research. The future of undergraduate research at the University of Toronto Mississauga is very bright indeed. I see Pareto growing into a well recognized and highly regarded journal for undergraduate research in Canada. I have personally edited and reviewed all of the articles in the first edition and I can attest to the quality of the research. I think you will enjoy the diversity of subject matter and the methodologies employed. I am so glad we are finally able to showcase the wonderful work and abilities of our economics students at the University of Toronto Mississauga. I truly believe we have some of the strongest students in the world. It has been my privilege and joy to advise and mentor many of them as they each begin their individual research journey.

I am indebted to those that helped make the first edition of Pareto possible. In particular I want to thank our Journal President Mary Kazek and VP marketing Nina (Zi Wei) Low. Mary has helped with edits and numerous other roles in getting the journal started. Nina has done an amazing job producing the branding and website design. Both have been on call for months helping wherever needed. The journal would not look half as good without them. I want to thank Khurshid Phirozmand (UEC President) for allowing us to embed Pareto into the UEC website and events calendar. I also give thanks to our economics chair, Margarida Duarte, who gave me the freedom and encouragement to start this new project. Finally I want to thank the student authors of the first edition: Gregory Chung, Ravdeep Sandal, Melissa Siqueira, Jun Takahashi and Renato Zimmermann. Their hard work and dedication to research should be commended and celebrated.

Dr. Nicholas Zammit
Assistant Professor, Teaching
Department of Economics
University of Toronto, Mississauga

Economics
UNIVERSITY OF TORONTO
MISSISSAUGA

Effects of Cannabis Legalization on Tobacco and Alcohol Prevalence

By Gregory Chung

This paper investigates the potential links between cannabis legalization and the use of other substances in the United States. In particular, it tests for changes in the prevalence of tobacco and alcohol use following cannabis legalization (Medical or recreational). This study is carried out on all 50 American states using a difference-in-differences estimator with data running from 2012 to 2018. It finds a significant decrease in tobacco use as a result of cannabis legalization. However, no statistically significant change in trend is observed in the prevalence of alcohol use following cannabis legalization.

***Keywords**: Cannabis use, drug policy, difference-in-differences, fixed effects analysis, panel data*

I. Introduction

Cannabis is the most widely used illicit drug in the USA (NSDUH statistics, 2018) with approximately 52 percent of all Americans aged 18 or older reporting having used it at least once (NBC News, 2017). The legalization of cannabis has been a hot topic during the past 10 years. Although cannabis use has remained illegal at the country/federal level since the 1930s in the United States, there are many individual states that have passed new legislation regarding the substance. This legislation ranges from the removal or reduction of criminal penalties (decriminalization) to complete legalization including commercial sales. As of April 2020, there are 39 states, including the District of Columbia, that have at least legalized the use of cannabis for medical purposes (see appendix figure A1 for breakdown). State legalization is on the rise for recreational cannabis as well [State Policies Department (2020)].

With these facts comes an important question: what is the relationship between the legalization of cannabis and the prevalence of tobacco and alcohol use? Relatively few studies have investigated this topic in the US. This research aims to add to the literature by attempting to evaluate the effects of the legalization of cannabis on tobacco and alcohol prevalence in the 50 states of the US and the District of Columbia. In particular, this paper will ask whether there is any complementarity or substitutability between cannabis and the use of tobacco or alcohol?

Being able to determine a relationship between cannabis and tobacco or alcohol could provide insights in developing policies relating to taxation, regulation and even pricing [Caulkins (2012)]. An early study conducted in Australia on cannabis, tobacco and alcohol, using data from the National Drug Strategy Household Surveys (NDSHS), found that cannabis and tobacco were complements [Cameron and Williams (2001)]. This paper aims to carry out a similar study using addiction data for the United States. This is relevant as it might inform policy regarding future legalization. If a link is found then policy targeted towards one substance might indirectly affect

demand for the other.

The rest of the paper will proceed as follows. Section 2 discusses the relevance of the topic and the extant literature, section 3 reviews data and methodology, section 4 presents the results, interpretation of coefficients and discusses the endogeneity issues and robustness of the model while section 5 concludes the paper.

II. The Extant Literature

There is a plethora of literature that aims to uncover the various implications of cannabis legalization. Studies range from estimating the health impacts of the substance to understanding the potential psychological and economic relationships cannabis has with other substances like alcohol and tobacco. In the latter case, investigations were carried out in an attempt to help decision makers devise proper regulatory policies to control a substance that was once illegal. These economic studies range from estimating the size of the black/illegal market to determining the price elasticity and degree of substitutability between legal and illegal cannabis [Amlung et al. (2019)]. Some studies have also attempted to estimate the substitutability with respect to alcohol and tobacco [Cameron and Williams (2001)].

Another relevant study was carried out in Switzerland to understand cannabis consumption modes amongst adolescents [Akre et al. (2009)]. A qualitative approach was adopted to identify cannabis and tobacco co-consumption and consumers perceptions of each substance. Five focus groups consisting of 22 youths (14 males) in the age group 15-21 years were interviewed seven times individually. Interestingly, this paper found that cannabis was perceived more positively than tobacco and therefore was considered a substitute to the latter. However, despite the perceived substitutability, the paper concluded that cannabis consumption could still induce nicotine dependence and cigarette smoking. It was thus recommended that the relationship between both substances was taken into consideration when implementing prevention programs.

A different conclusion was reached by [Cameron and Williams (2001)]. This study used individual level data for the years 1988, 1991, 1993 and 1995 from the National Drug Strategy Household Survey alongside alcohol and tobacco price indices from the Australian Bureau of Statistics. The data was pooled, leading to a sample size of 9744 observations. The study concluded that cannabis and tobacco were complements, while alcohol and cannabis were substitutes. A key aspect of the study was that participation in all 3 drugs was sensitive to own price changes and that decriminalization of cannabis resulted in an increase in its consumption.

One possible factor that could account for the discrepancy in the findings is related to the type of data collected. As mentioned by [Verbic et al. (2001)], the majority of studies use general population surveys for estimating the prevalence of cannabis. Some authors have argued that this type of data has a drawback of not being able to reach the cannabis-using population because of users unwillingness to report their true conditions due to fear of stigmatization [Fendrich and Johnson (2005)]. Other researchers favour the use of information collected from the household population surveys in understanding marijuana markets [Caulkins and Pacula (2016)].

The lack of reliable price data for cannabis also hampers empirical estimation of price responsiveness. Economists have tried to estimate price elasticity wherever price data were available. Since available prices are not always idiosyncratic, estimated price effects differ broadly across studies. Due to the inexistence of data on quantities consumed, price elasticities are generally

reported in terms of participation probability (as cited in Verbic et al, 2019).

A recent study, conducted by [Veligati et al. (2020)] in the US, investigated the impact of the legalization of medical and recreational cannabis on state-level per capita alcohol and cigarette consumption. The authors used data from state tax receipts maintained by the Centers for Disease Control and National Institute for Alcohol Abuse and Alcoholism with a difference-in-differences estimator. One strength of this paper was the use of a 3-tiered model that included different covariates to mimic the differences in medical and recreational legislations across states. The authors found no significant relationships between medical or recreational cannabis policies and per capita sales of cigarettes and alcohol.

While most economic studies on the topic used data on economic variables (such as prices, sales and taxes), relatively few of them used data on the actual number of people consuming the substances. Therefore, this paper aims to carry out a similar study to [Veligati et al. (2020)] with a difference-in-differences method, but using a drugs prevalence data set for the US.

III. Data & Methods

A. Data Summary

Following [Caulkins and Pacula (2016)] this paper focuses on household substance use. Specifically, this paper employs household data obtained from the Substance Abuse and Mental Health Service Administration (SAMHSA) website, which operates under the US Department of Health Human services. The data set consists of seven yearly iterations of the National Survey on Drug Abuse and Health (NSDUH) data over the period 2012 to 2018 for all of the 51 states. The data were collected from a sample of individuals who have a tendency towards addiction. Tobacco prevalence in our data set is defined as the number of people (in thousands) having reported consuming cigarettes, smokeless tobacco (i.e., snuff, dip, chewing tobacco, or "snus"), cigars, or pipe tobacco for an amount of at least 100 cigarettes in their lifetime. Alcohol prevalence includes the number of people (in thousands) having reported consuming at least a full drink in the past 30 days. Given the approach taken in this paper, the age group 26 and older will be the group of interest. Table 1 provides a summary of the data for the age category of interest.

B. Survey Breakdown

Data used in this study was obtained from an annual household interview survey of substance use, substance use disorders, mental health and the receipt of treatment services for the disorders (see appendix figure A2 for the survey front page). The survey is targeted towards non-institutionalized individuals of age 12 years and older. In particular, the survey measures the average prevalence of drug and substance use at the state level for the following age groups: 12-18, 19-25 and 26 and above. Responses from approximately 67500 households are recorded annually. The survey is a cross-sectional one, carried out each year (on different households).

An interesting trend can be immediately observed in these statistics. The mean cannabis prevalence is lower for states where the substance is legal than for states where it is illegal. A potential explanation for this observation is the control of quantity and hoarding effect. When cannabis is illegal, (the supply of the substance is very low), people tend to consume more of it and stock it up. Conversely, in states where the supply is higher (coming from the legalization),

people no longer have the need to hoard. Another reason might be the "criminality effect" as put forward by Becker, Murphy and Grossman in their paper "The Economic Theory of Illegal Goods: The Case of Drugs" [Becker et al. (2006)]. According to Becker, legal taxation on producers may be more effective than enforcement of illegal drugs in reducing consumption. The survey data does reveal higher mean for first time use of cannabis for legal states. Intuitively, once cannabis is available through various point of sales (POS), it is easier to access leading to greater first time use. However, the mean use of cannabis overall is not significantly different from legal to illegal states. Difference of means testing for each type of substance use (cannabis, alcohol and tobacco) does not reveal a statistically significant difference amongst states where cannabis is legal vs. ones where it is illegal. Hence, the descriptive statistics do not provide a clear picture on the causality between prevalence and legalization but signal some potential links (for example with average first use being much higher in states that have legalized).

Table 1—Summary statistics on prevalence (000s) for the group 26 years and older

	Illegal (N=84)				**Legal (N=273)**			
	Mean	sd	Min	Max	Mean	sd	Min	Max
Cannabis	269.869	221.717	13	829	265.205	364.052	17	2676
Alcohol	2357.464	2210.118	197	8941	2219.033	2593.35	253	14513
Tobacco	1181.988	994.183	103	4189	963.267	963.772	101	4538
First time	4.036	3.254	0	15	8.381	11.089	0	96

C. Methodology

The repeat cross-sectional surveys do not allow for estimation of within-individual household changes in outcomes over time. Analysis therefore must focus on estimating outcomes averaged across states, and in comparing changes in these over time between 2012 and 2018. This paper assumes that determinants other than the legalization of cannabis remained stable in the fifty states over time or followed a parallel change. Given this assumption, a difference-in-differences (DID) analysis can be used to uncover the average net effect of the legalization on tobacco prevalence amongst individuals age 26 or above.

Basically, the DID estimator can be calculated as:

$$y_{\text{DID}} = (y^{\text{L}}_{\text{AL}} - y^{\text{L}}_{\text{BL}}) - (y^{\text{I}}_{\text{AL}} - y^{\text{I}}_{\text{BL}})$$

Where superscript L and I refer to the states where cannabis is legal and illegal respectively and subscript AL is the time period after the legalization while BL is before legalization.

The use of Pooled OLS (POLS) to calculate the DID estimator is justified because of the repeated cross-sectional nature of the survey. The sampling technique was designed so that no overlap amongst residents was expected. The only way overlap could happen was if the individuals moved to another area segment/state in-between years and their new residence was selected again the following year for the survey. By design, panel data techniques are not

necessarily required. A Breusch-Pagan test was used to identify any issues and help with model selection. According to the test, the null hypothesis is constant variance (homoskedasticity) within the sample. The test results indicate there was no statistical evidence at the 5% level to reject the null. For the main approach in this study, panel data methods were discarded and the POLS model was selected.

Although the Breusch-Pagan test did not detect any issue of heteroskedasticity in the sample, the model was also estimated on the full sample using a panel data treatment to account for unobserved heterogeneity and test for the robustness of the POLS estimation. The Hausman test was used to determine between either a random effects or a fixed effects estimation technique. Recall that the null hypothesis for the test is that a RE model is consistent and preferred. Test results indicated that there was evidence to reject the null, so that a FE model was chosen, in addition to the main model estimated. Test results relating to model selection are available on request.

The basic DID coefficient (Y_DID) is estimated by the following equation using Pooled Ordinary Least Squares:

$$\begin{aligned} ln\text{Tobacco}_{\text{it}} = \beta_0 + \gamma\text{Legal}_{\text{i}} + \theta\text{AfterYear}_{\text{t}} + \delta(\text{Legal} * \text{AfterYear}_{\text{it}}) + \\ \eta\text{Recreational}_{\text{i}} + \phi\text{Decriminalized}_{\text{i}} + \varepsilon \end{aligned}$$

The outcome variable $ln\text{Tobacco}_{\text{it}}$ is the natural logarithm of tobacco prevalence, Legal_{i} is a dummy for the treatment (legalization) effect, $\text{AfterYear}_{\text{t}}$ is a time dummy to indicate whether an observation occurs after the legalization and (Legal $*$ AfterYear) is an interaction term that gives the average treatment effect. Before going any further, it is important to define what is meant by legalization here. In the regressions, 'Legal' is a dummy variable that refers to cannabis being legal for either medical or recreational use. In fact, if a state legalizes the substance for recreational purposes, it automatically means that it is legal for medical use. A recreational use dummy is used to capture the effect of a full legalization for recreational purposes. Many states have ideosyncratic policies with regard to the legality of cannabis. There is a spectrum of legislations between illegal and legal. To account for this, we also include a state dummy for decriminalized (possession of cannabis not seen as a criminal offence).

For the sake of specification and statistical relevance, it is also important to account for both state and year fixed effects. Economically, having state fixed effects helps to control for the possible unobserved heterogeneity such as varying state legislations that affects prevalence/use of the drug amongst households within a particular state. Year fixed effects are expected to capture unobserved heterogeneity in a particular year. λ_{i} captures the state fixed effects while μ_{i} captures time fixed effects. One advantage of the DID method is that it mimics an experimental research design by designating a treatment and control group using observational data. The model can also allow for dummies/covariates to be included to account for the different directions and characteristics the two groups can take.

The adjusted fixed effects Pooled OLS is now:

$$\begin{aligned} ln\text{Tobacco}_{\text{it}} = \beta_0 + \gamma\text{Legal}_{\text{i}} + \theta\text{AfterYear}_{\text{t}} + \delta(\text{Legal} * \text{AfterYear}_{\text{it}}) + \\ \eta\text{Recreational}_{\text{i}} + \phi\text{Decriminalized}_{\text{i}} + \lambda_{\text{i}} + \mu_{\text{t}} + \varepsilon \end{aligned}$$

Where the treatment group is defined as states that have legalized cannabis prior to 2018 and the control group consists of those states where cannabis is still illegal. To evaluate the specification of the model after the inclusion of these additional covariates, a Ramsey Regression Equation Specification Error Test (RESET) was used. The test found no evidence at the 5% level to reject the null and concluded that the model is well specified.

IV. Results

A. Full Sample Specification

The results of the regressions are shown in Table 2. To facilitate the interpretation of coefficients, we make use of a log-level model. The coefficients of the dependent variables are interpreted as a $100 * (e^{\beta} - 1)$ percentage change in average tobacco prevalence. Keeping all else constant, being a legal state increases tobacco prevalence by approximately 1.37%. This coefficient is not significant. The coefficient of the dummy for recreational legalization says that the average tobacco prevalence decreases by 52.3% if a state fully legalizes cannabis for recreational purposes. Conversely, states where the recreational use of cannabis is decriminalized (not a criminal offence) are likely to experience a rise of 5.8% in the average prevalence of Tobacco. Both coefficients are statistically significant at the 1% level.

Our main coefficient of interest, DID (which is the average net treatment effect), is calculated by taking an unweighted average of all the DID coefficients from the years 2012 to 2017. The resulting average is -0.0084482. Given individual significance tests for each coefficient we can conclude that the average DID effect is significant at the 5% level. This means that on average over the 6 years (excluding 2018 because of no post 2018 data), average tobacco prevalence fell by 0.845% as a result of the legalization of cannabis, for either medical or recreational use.

Next we apply the same model using average alcohol prevalence as the dependent variable. We find that when a state is legal, there is an approximate increase of 73.3% in the average alcohol prevalence and this is significant at the 1% level. Similar to the estimations on tobacco, decriminalizing cannabis has a positive effect on the average Alcohol prevalence. In fact, there is an approximate increase of 93.9%. The coefficient on the recreational legalization measure points towards a 79.8% fall in alcohol prevalence as a result of legalizing cannabis for recreational use and is statistically significant at the 1% level.

Following the same averaging process in the case of tobacco prevalence, we obtain that the average net treatment effect over the time period 2012 to 2018 is a decrease of 0.3% in alcohol prevalence. Unlike the calculated average from the tobacco regression, it is not statistically significant for alcohol. The reported R-squared values are 0.99 for both the tobacco and alcohol regressions, indicating a very good fit but potentially the presence of some specification issues or multicollinearity.

These results seem to point towards a substitution away from tobacco products resulting from the legalization of cannabis for either medical or recreational use. However, the same causal relationship cannot be established for alcohol.

Table 2—Results: Full sample

	(1) Tobacco	(2) Alcohol
Legal	0.0137	0.553***
	(0.0203)	(0.0194)
DID12	0.00614	0.00163
	(0.0213)	(0.0181)
DID13	-0.0128	-0.00760
	(0.0196)	(0.0153)
DID14	0.00381	-0.00653
	(0.0196)	(0.0123)
DID15	0.0121	-0.0111
	(0.0210)	(0.0138)
DID16	-0.0248	-0.00261
	(0.0225)	(0.0159)
DID17	-0.0352	0.00819
	(0.0234)	(0.0173)
Recreational	-0.741***	-1.599***
	(0.0188)	(0.0146)
decriminalized	0.0580***	0.662***
	(0.0217)	(0.0216)
N	357	357
R^2	0.998	0.999

Robust standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

As a final check on these results, a panel data model using FEs was estimated for the full sample. The coefficients do not change significantly and the average DID effect is still -0.00845 and is significant at the 5% level. The model did not suffer from heteroskedasticity. Table IV.A reports the detailed results of the regressions including FEs for the full sample.

Table 3—Results: Full sample Panel regression

	(1) Tobacco	(2) Alcohol
Legal	0	0
	(.)	(.)
DID12	0.00614	0.00163
	(0.0211)	(0.0148)
DID13	-0.0128	-0.00760
	(0.0211)	(0.0148)
DID14	0.00381	-0.00653
	(0.0211)	(0.0148)
DID15	0.0121	-0.0111
	(0.0211)	(0.0148)
DID16	-0.0248	-0.00261
	(0.0211)	(0.0148)
DID17	-0.0352*	0.00819
	(0.0211)	(0.0148)
Recreational	0	0
	(.)	(.)
Decriminalized	0	0
	(.)	(.)
Year and state fixed effects	Yes	Yes
N	357	357
r2_within	0.257	0.292
r2_between	0.00120	0.000146
r2_overall	0.000931	0.000238

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B. Restricted Sample Specification

To cross-check the reliability of the main methodology and evaluate the statistical significance of the DID coefficient, an alternative approach was also adopted. This method consists of pooling states into subsamples based on their respective periods of legalization. Specifically, this approach uses restricted samples of states, for each legalization year between 2012 and 2018, while excluding 2018 (no post-2018 data). The subsamples are made up of the states that legalized either medical or recreational cannabis in the years 2012 to 2017 and 12 states where cannabis remained illegal. The largest subsample is that of legalization year 2014, where the treatment group is made up of 7 states that passed new legislation in 2014. The subsample for 2016 is the second largest with 4 states in the treatment group. Subsamples 2012, 2015 and 2017 all have 3 states each in the treatment group. Finally, subsample 2013 has only 2 states that implemented new cannabis legislation.

The reported coefficients (Table 3) for 'Legal' are negative for the subsample of states with legalization year 2014, 2015 and 2016, and are all statistically significant at the 1% level. This indicates a switch away from tobacco in states that legalized cannabis for either medical or recreational use. However, this paper is mostly interested in the respective DID coefficients. It is worth noting that only the DID coefficients from legalization years 2012, 2015 and 2016 are statistically significant and represent a fall of 5.91%, 4.73% and 5.26% in the average tobacco prevalence respectively. The only positive coefficient for the interaction term is recorded in 2014 only and is insignificant.

The same approach is applied with alcohol as the dependent variable. Table 4 reports the estimated coefficients. Notice here that, once again, the only significant DID coefficients are negative. These are recorded for the years 2013, 2014 and 2015 respectively. An interesting observation is that in 2015, after legalization there is a fall in both tobacco and alcohol prevalence. This subsample included the states of Georgia, Louisiana and Texas. All three of them legalized cannabis for medical use only. Georgia and Texas had similar legislation, whereby CBD oil containing less than 5% and 0.5% THC was legalized. Hence, these were very unique legalizations and could explain outliers in the results.

This alternative approach validates what was found in the first method in the case of tobacco prevalence. There has been a substitution away from tobacco that resulted from the legalization of cannabis. The results are less clear for alcohol. Although the alternative method captured significant downward trends resulting from the legalization, no statistically significant evidence was found using the main approach. Therefore, it is difficult to claim a substitution away from alcohol.

Table 4—Restricted samples approach: Tobacco

	(1) Tobacco	(2) Tobacco	(3) Tobacco	(4) Tobacco	(5) Tobacco	(6) Tobacco
Legal	0.725***	0.170***	-2.213***	-1.196***	-1.762***	0.00695
	(0.0322)	(0.0402)	(0.0206)	(0.0310)	(0.0256)	(0.0165)
Recreational	-0.0585***	-0.0585***	-0.0585***	-0.0585***	-0.0585***	-0.0585***
	(0.0213)	(0.0196)	(0.0206)	(0.0211)	(0.0204)	(0.0201)
Decriminalized	0.0580**	0.0580**	0.0580**	0.0580**	0.0580**	0.0580**
	(0.0239)	(0.0230)	(0.0264)	(0.0238)	(0.0238)	(0.0240)
DID12	-0.0591*					
	(0.0311)					
DID13		-0.0207				
		(0.0314)				
DID14			0.00923			
			(0.0162)			
DID15				-0.0473**		
				(0.0233)		
DID16					-0.0526**	
					(0.0234)	
DID17						-0.0129
						(0.0302)
N	105	98	133	105	112	105
R^2	0.998	0.998	0.998	0.998	0.998	0.998

Robust Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5—Restricted samples approach: Alcohol

	(1) Alcohol	(2) Alcohol	(3) Alcohol	(4) Alcohol	(5) Alcohol	(6) Alcohol
Legal	0.886***	0.272***	-1.662***	-0.940***	-1.174***	0.541***
	(0.0260)	(0.0217)	(0.0166)	(0.0193)	(0.0206)	(0.0160)
Recreational	0.0766***	0.0766***	0.0766***	0.0766***	0.0766***	0.0766***
	(0.0155)	(0.0157)	(0.0157)	(0.0157)	(0.0153)	(0.0157)
Decriminalized	0.662***	0.662***	0.662***	0.662***	0.662***	0.662***
	(0.0216)	(0.0217)	(0.0210)	(0.0214)	(0.0214)	(0.0214)
DID12	-0.01000					
	(0.0282)					
DID13		-0.0667***				
		(0.0174)				
DID14			-0.0221*			
			(0.0123)			
DID15				-0.0471***		
				(0.0177)		
DID16					-0.0206	
					(0.0159)	
DID17						0.0000196
						(0.0251)
N	105	98	133	105	112	105
R^2	0.999	0.999	0.999	0.999	0.999	0.999

Robust Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

V. Conclusions

One should be careful when interpreting the results of this study. Although we found a statistically significant fall in the average tobacco prevalence resulting from new legislation on marijuana, it does not necessarily mean that both substances are economic substitutes. As mentioned earlier, different individuals have different preferences with respect to drugs. Some might be more prone to consuming cannabis because they are already consuming tobacco and vice versa. More sophisticated structural models can solve for that. Our study, on the other side,

focuses more on changes in trends in average prevalence. What makes this study novel is the use of the unique NSDUH data set to account for the potential trends in substance prevalence among individuals who have a tendency towards addiction. However, this can be a potential source of bias in our estimates, since we are tracking individuals who are prone to substance and drug addiction. Another weakness of this study is the lack of fine-grained data/granulated data. In fact, this data set is limited to the state and year level, so that the detailed effects of cannabis legislations cannot be clearly captured. Even state and year fixed effects cannot fully capture the effects of such varying legislations. Therefore, having access to the full data set with state identifiers would produce more reliable and accurate estimations. Despite this the research presented in this paper offers a glimpse into the potential links between cannabis legalization and the use of tobacco and alcohol. Even with limited data its clear there are strong correlations between the use of each substance. This has significant implications for policy makers that are considering legalization.

Overall, although limited by the absence of a more granulated data set, this paper finds a negative and significant relationship between average tobacco prevalence and the legalization of cannabis for either medical or recreational use. It finds no strong positive or negative relationship between alcohol and cannabis. These findings do not parallel exactly the conclusions reached by some similar studies that found no significant association between the legalization of cannabis and sales of cigarettes and alcohol [Veligati et al. (2020)]. While alcohol use seems unaffected by legalization of cannabis the coefficient estimates relating tobacco use to cannabis legalization are large and significant. Given the desire of politicians to limit tobacco use, these results could play a major role in considering future policy decisions around cannabis legalization.

The question of substitutability or complementarity between cannabis and alcohol and tobacco does not have a 'one-size-fits-all' type of answer [Guttmannova et al. (2016)]. In fact, various studies carried out in the past on the subject had contrasting results. One potential explanation for that can be found in the psychology literature, whereby an individual pattern of consumption for drugs depends on preferences and motivations. These preferences vary widely across individuals [Simons et al. (2005)]. Future studies should aim towards obtaining more granulated individual level data and the inclusion of more covariates such as age, gender, and even taxes on alcohol, tobacco and cannabis consumption. The use of more sophisticated structural models can also be used to account for the endogeneity problem between the three substances to shed more light on these notoriously-debated relationships.

REFERENCES

[Akre et al. 2009] Akre, C. ; Michaud, P. ; Berchtold, A. ; Suris, J.: Cannabis and tobacco use: where are the boundaries? A qualitative study on cannabis consumption modes among adolescents. In: *Health Education Research* 25 (2009), Nr. 1, S. 74–82

[Amlung et al. 2019] Amlung, M. ; Reed, D. ; Morris, V. ; Aston, E. ; Metrik, J. ; MacKillop, J.: Price elasticity of illegal versus legal cannabis: a behavioral economic substitutability analysis. In: *Addiction* 114 (2019), Nr. 1, S. 112–118

[Becker et al. 2006] Becker, G. ; Murphy, K. ; Grossman, M.: The Market for Illegal Goods: The Case of Drugs. In: *Journal of Political Economy* 114 (2006), Nr. 1, S. 38–60

[Cameron and Williams 2001] Cameron, L. ; Williams, J.: Cannabis, Alcohol and Cigarettes: Substitutes or Complements? In: *Economic Record* 77 (2001), Nr. 236, S. 19–34

[Caulkins 2012] Caulkins, J.: Design considerations for legalizing cannabis: lessons inspired by analysis of California's Proposition 19. In: *Addiction* 107 (2012), Nr. 5, S. 865–871

[Caulkins and Pacula 2016] Caulkins, J. ; Pacula, R.: Marijuana Markets: Inferences from Reports by the Household Population. In: *Journal of Drug Issues* 36 (2016), Nr. 1, S. 173–200

[Fendrich and Johnson 2005] Fendrich, M. ; Johnson, T: Race/ethnicity differences in the validity of self-reported drug use: results from a household survey. In: *Journal of Urban Health* 82 (2005), Nr. 2, S. 67–81

[Guttmannova et al. 2016] Guttmannova, K. ; Lee, C. ; Kilmer, J. ; Fleming, C. ; Rhew, I. ; Kosterman, R. ; Larimer, M.: Impacts of Changing Marijuana Policies on Alcohol Use in the United States. In: *Alcohol: Clinical and Experimental Research* 40 (2016), Nr. 1, S. 33–46

[Simons et al. 2005] Simons, J. ; Gaher, R. ; Correia, C. ; Hansen, C. ; Christopher, M.: An affective-motivational model of marijuana and alcohol problems among college students. In: *Psychology of Addictive Behaviors* 19 (2005), Nr. 3, S. 326–334

[State Policies Department 2020] State Policies Department, MPP: *Marijuana Policy Progress Report 2019 Legislative Update*. 2020. – URL `https://www.mpp.org/issues/legislation/`

[Veligati et al. 2020] Veligati, S. ; Howdeshell, S. ; Beeler-Stinn, S. ; Lingam, D. ; Allen, P. ; Chen, L. ; Grucza, R: Changes in Alcohol and Cigarette Consumption in Response to Medical and Recreational Cannabis Legalization: Evidence from U.S. State Tax Receipt Data. In: *International Journal of Drug Policy* 102585 (2020)

[Verbic et al. 2001] Verbic, M. ; Cok, M. ; Peric, P.: An economic analysis of the illegal marijuana market in Slovenia. In: *Economic Research* 32 (2001), Nr. 1, S. 657–672

APPENDIX

Order	State	Legalization year	Order	State	Legalization year	Order	State	Legalization year
1	Alabama	N/A	20	Maine	1999	40	Rhode Isla	2006
2	Alaska	2014	21	Maryland	2013	41	South Car	2014
3	Arizona	2010	22	Massachu	2012	42	South Dak	N/A
4	Arkansas	2016	23	Michigan	2008	43	Tennessee	N/A
5	California	1996	24	Minnesota	2014	44	Texas	2015
6	Colorado	2012	25	Mississipp	N/A	45	Utah	2014
7	Connecticu	2012	26	Missouri	2018	46	Vermont	2004
8	Delaware	2011	27	Montana	2004	47	Virginia	N/A
9	District of	1998	28	Nebraska	N/A	48	Washingto	1998
10	Florida	2017	29	Nevada	2000	49	West Virgi	2017
11	Georgia	2015	30	New Hamp	2013	50	Wisconsin	N/A
12	Hawaii	2000	31	New Jerse	2010	51	Wyoming	N/A
13	Idaho	N/A	32	New Mexic	2007			
14	Illinois	2014	33	New York	2014			
15	Indiana	2017	34	North Caro	N/A			
16	Iowa	2014	35	North Dak	2016			
17	Kansas	N/A	36	Ohio	2016			
18	Kentucky	N/A	37	Oklahoma	2018			
19	Louisiana	2015	38	Oregon	1998			
20	Maine	1999	39	Pennsylva	2016			

Figure A1. List of states with legalization dates

INTRODUCTION AND INFORMED CONSENT FOR INTERVIEW RESPONDENTS AGE 18+

INTRODUCE YOURSELF AND STUDY AS NECESSARY: Hello, I'm _____, and I'm working on a nationwide study sponsored by the U.S. Department of Health and Human Services. You should have received a letter about this study. (SHOW LEAD LETTER, IF NECESSARY.)

READ THE BOXED INFORMATION BELOW BEFORE STARTING EVERY INTERVIEW

This year, we are interviewing about 70,000 people across the nation. You have been randomly chosen to take part. You will represent over 4,500 other people who are similar to you. You may choose not to take part in this study, but no one else can take your place. We will give you $30 when you finish the interview.

GIVE STUDY DESCRIPTION TO R IF YOU HAVE NOT ALREADY DONE SO.

This study asks about tobacco, alcohol, and drug use or non-use, knowledge and attitudes about drugs, mental health, and other health issues. It takes about an hour. You will answer most of the questions on the computer, so I will not see your answers. We are only interested in the combined responses from all 70,000 people, not just one person's answers. This is why we do not ask for your name and we keep your answers separate from your address. RTI may contact you by phone or mail to ask a few questions about the quality of my work. This is why we ask for your phone number and current address at the end of the interview.

While the interview has some personal questions, federal law keeps your answers private. We hope that protecting your privacy will help you to give accurate answers. You can quit the interview at any time and you can refuse to answer any questions.

If it is all right with you, let's get started.

(Can we find a private place to complete the interview?)

FIGURE A2. NSDUH SURVEY FRONT PAGE

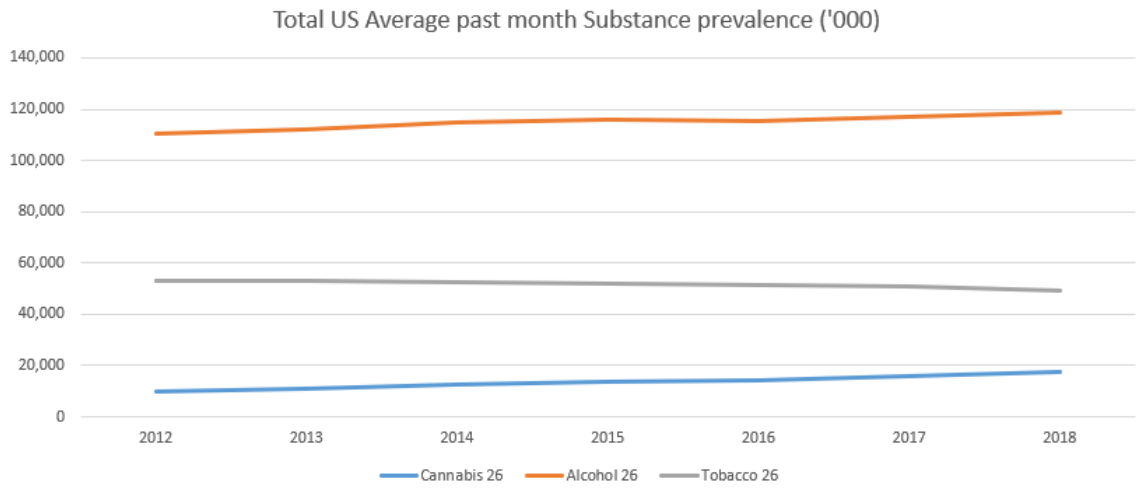

Figure A3. Trends US total

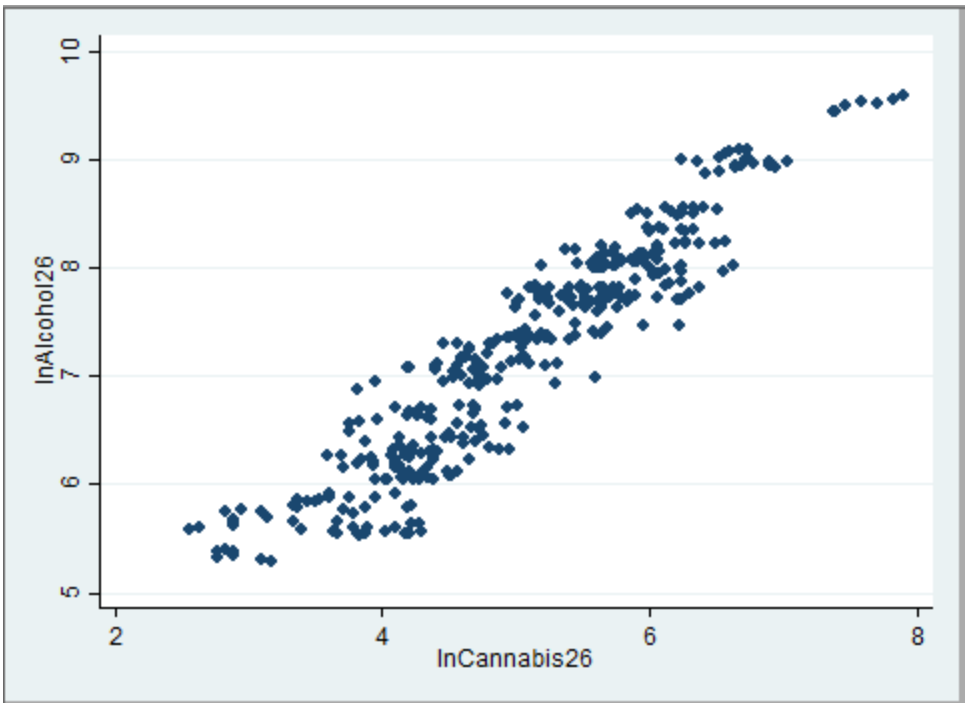

Figure A4. Correlation log Tobacco and log Alcohol

Skilled Immigration and its Impact on Ontario's Unemployment Rates

By Ravdeep Sandal

This study exploits policy differences across Canadian provinces to analyze the effects of skilled immigration on the unemployment rate. In 2005, Ontario signed its first Provincial Nominee Program agreement. In 2007, Ontario adopted and began to utilize the Provincial Nominee Program, a program designed to fill labour shortages and meet economic demands. To analyze the impact of such a policy, Ontario is compared to Quebec, which has not adopted the Provincial Nominee Program. This report utilizes ARIMAX models to find the impact of this skilled labour program on unemployment levels. The results demonstrate that the introduction of the Provincial Nominee Program in Ontario has led to a statistically significant decline in the unemployment rate. Hence, this inflow of skilled labour provides a net benefit, which has implications for more liberal policies regarding skilled migration.

***Keywords**: Immigration, Labour Economics, Unemployment, ARIMAX, Provincial Nominee Program*

I. Introduction

Canada is a country with its foundations built on immigration. By 2011, individuals in Canada classed as immigrants made up 20.7% of the population. Those classed as immigrants are anticipated to constitute roughly 30% of the population by 2036 [Canada (2022)]. With 431,645 new permanent residents to Canada in 2022 and an increased quota of 447,055 new permanent residents in 2023, immigration continues to play a vital role in sustaining Canada's labour force and population growth [Canada (2023)]. Approximately 75% of Canada's population growth is due to immigration, with most of this comprised mainly of economic migrants, i.e. skilled immigration.

In Canada immigration is vital in combating a shrinking working-age population with potential labour shortages [Latif (2015)]. To illustrate, the majority of the health care sector is comprised of immigrant workers. Immigrant labour represents 23% of registered nurses, 37% of pharmacists, 39% of dentists, 36% of physicians, and 35% of nurse aides [Canada (2023)]. Immigration has provided Canada with a continuous influx of about 200,000 immigrants annually since the late 1950s. Canadian immigration levels have steadily increased, recently doubling in size. The Trudeau Liberal government continues to demonstrate an interest in increasing immigration quotas [El-Assal and Fields (2018)].

In 1967, Canada introduced its points-based immigration system. Since then, Canada has continued to refine its immigration system, introducing the Express Entry system (EE) in 2015. Over the years, Canada has reduced its quotas for non-economic immigration streams while

increasing economic immigration. In 2023, about 57% of new permanent residents will be from the economic stream. Canada has become a nation that prioritizes skilled labour above all else. With such high levels of immigration, the impact on employment of native workers comes into question. Domestic workers must now compete with an influx of skilled competitors, possibly resulting in an increase in native unemployment.

Previous literature has found that as immigration-reliant countries experience pockets of unemployment, economic growth stagnates and declines [Islam (2007)]. The fear of stagnating growth has introduced new concerns regarding immigration in developed countries, including Canada. Some Canadians believe foreign nationals may steal jobs from native Canadian workers. Others believe increased immigration puts pressure on taxpayers when immigrants struggle to find employment due to skill limitations. These Canadians attribute unemployment to increased immigration; however, others believe immigration reduces skill shortages in the economy, which the Canadian government desires [Islam (2007)].

The impact of immigration on growth can be ambiguous, being either negative or positive. The negatives associated with increased immigration may be attributed to reduced employment opportunities for native workers. Low-skilled and noncompetitive native workers may quickly find new immigrants becoming strong competitors, increasing the unemployment rate for native workers. Immigrants who displace workers (or immigrants who remain unemployed) will increase overall unemployment, increasing stresses on the tax system. Native worker bargaining power may decline as competing and equally qualified immigrants may be willing to accept lower wages. As a result, native worker wages may decrease [Islam (2007)].

However, it is possible for increased immigration to benefit both foreign and native workers. New immigrants can directly increase the demand for goods and services, developing new job opportunities and aiding native workers. Aggregate expenditure will also experience positive shocks, directly through increased immigrant spending and indirectly through industrial and government expenditures. The increased expenditure can lower unemployment by creating new jobs and improving wages [Islam (2007)]. Immigrants have weaker bargaining positions, resulting in lower wages than native workers. With an abundance of immigrant workers, a decline in the average wage can lead to an increase in the demand for labour as labour costs decrease. This decline may directly benefit native workers as an increase in jobs will result in reduced unemployment. The new employment surplus can increase native workers' bargaining positions, allowing native workers to bargain for higher wages [Islam (2007)]. Therefore, immigration can benefit both parties: immigrants and native workers. Given there exists conflicting theoretical equilibria, this paper will use an empirical approach to analyze the effect of skilled immigration on Ontario's unemployment rate.

II. Literature Review

The discussion of immigration and its impact on the economy can be sensitive and often becomes topical during economic crises. During the 2008 global recession, many politicians debated the consequences of immigration. The fear of immigrants stealing native workers' jobs became a central issue [Rohac (2014)]. Although Canadians generally hold a positive view on immigration, some believe immigrants threaten native workers [Challinor (2011)]. In the past, most studies focused on theoretical analysis and hypothetical situations; however, recent studies

have shifted towards empirical evidence. Modern research utilizes panel data and time-series modelling to analyze the implications of immigration [Latif (2015)].

As mentioned in the introduction, previous literature discussing the causal impact of immigration on unemployment has not arrived at a unanimous conclusion. One of the earliest theoretical analyses, conducted by Berry and Soligo, on the welfare implications of immigration found a positive link [Berry and Soligo (1969)]. Their research was in favour of immigration as it revealed that increased immigration led to the improvement of the economic situation of native workers. Winter-Ebmer and Zweimuller also found support for the positive relationship between immigration and growth; their study found the number of immigrant workers positively correlates with native worker earnings [Winter-Ebmer and Zweimller (1999)]. More recent studies, such as Ortega's, have looked at theoretical migration and non-migration equilibria, suggesting that increased migration will create new jobs in host countries [Ortega and Peri (2014)]. These equilibria suggest that increasing immigration leads to a decline in unemployment and an increase in local wages.

Contrary to these positive studies, many researchers have found negative effects resulting from increased immigration; while other studies have found no impact at all. Harris and Todaro utilized a two-sector model to analyze unemployment and migration, finding that immigration negatively affects native workers [Harris and Todaro (1970)]. Other researchers found no significant impact of immigration on native job opportunities. Some include [Gianmarco and Giovanni (2012)], [Card (2001)], [Card (2005)], [Card (2007)] and [Borjas (2006)].

This paper will contribute to the existing literature as it builds on previous research and brings a new approach to the analysis of skilled immigration on unemployment. This paper will focus on the effects of introducing a new skilled-immigration program, the 2007 Ontario Provincial Nominee Program (PNP). The system was initially introduced in 1996 and its first implementation was in 1999 [Canada (2011)]. Ontario was a late adopter of the system in 2007. Central Canada was a leading region in adopting new immigration policy; however, Ontario and Quebec have had very different immigration policies and processes within the region. The federal immigration process is the same for all provinces and territories, but each province and territory can design its own PNP immigration process. The PNP system operates identically across Canada, but applicant requirements differ according to provincial needs and goals. All provinces and territories utilize their PNP system based on occupational demands; however, Quebec is the only province that does not use the PNP system. By comparing the incorporation of the PNP system in Ontario to to its absence in Quebec, this paper will determine how the new system has impacted provincial unemployment.

III. Summary of the Canadian Immigration System

The Canadian immigration system has three main components: economic immigration, family reunification, and humanitarian grounds.

The first component, the economic immigration stream, is the primary immigration stream into Canada. This stream is centred around a points-based system, ranking individuals according to their ability to adapt to the Canadian work environment. Individuals are awarded points based on their education, experience, skills, age, and other adaptability factors. This economic immigration stream branches into two categories: Express Entry (EE) and Provincial Nominee

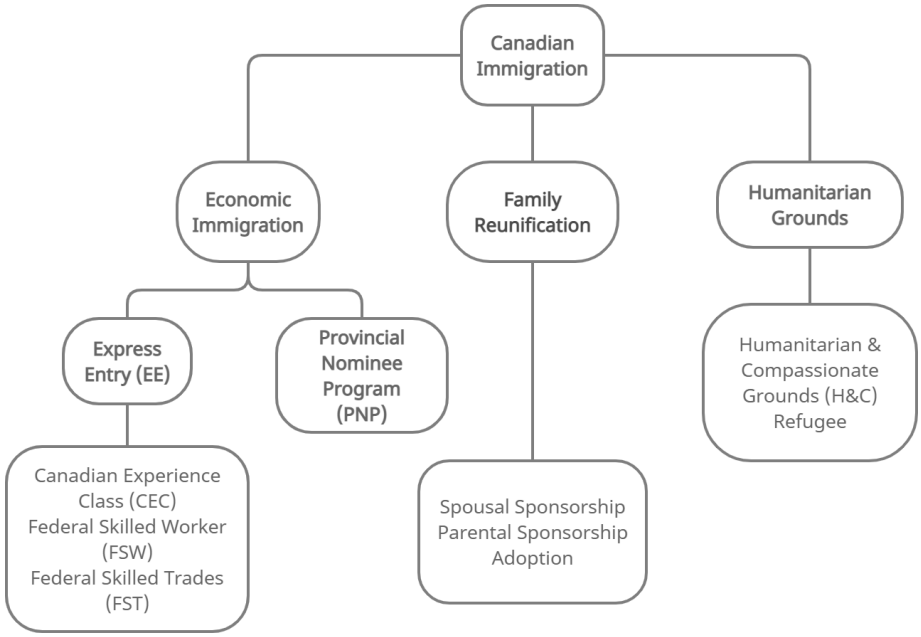

Figure 1. Visual Representation of the Canadian Immigration System

Programs (PNP). The EE program is a federal program in which individuals may immigrate to any province or territory. To qualify, immigrants must meet a points threshold and the requirements of one of the underlying categories: Canadian Experience Class (CEC), Federal Skilled Worker (FSW), or Federal Skilled Trade (FST).

The PNP program is province and territory-specific, and individuals may immigrate to the province or territory that nominates them for immigration (Quebec does not partake in the PNP system). The PNP system allows provinces to nominate foreign individuals for immigration to Canada based on provincial/territorial occupational demands, workforce shortages and development/advancement plans. These immigrants must meet language requirements and meet province and territory specific thresholds. The core aspect of the PNP system is to utilize highly skilled labour to meet labour shortages, strengthen local economies and kickstart new industries [Canada (2011)].

The second immigration stream, family reunification, focuses on reuniting individuals with family members. This stream focuses on sponsorship programs, with its core entry for parental/spousal sponsorship and adoption.

The third immigration stream, humanitarian grounds, is designed to provide access for displaced individuals and those at risk. This stream also allows illegal citizens in Canada to achieve permanent residence based on their establishment and ties to Canada and for other humanitarian reasons.

IV. Data

Data in this study comprises multiple sources, primarily from Statistics Canada. Federal-level data was collected from Statistics Canada and provincial-level data was sourced from provincial and territorial data banks. Data on the gross domestic product, labour force, minimum wage, migration (immigration and emigration), consumer price index, population, and Canada wide unemployment rates are readily available from Statistics Canada. Provincial-level unemployment rates are available on The Alberta Economic Dashboard. Recession and Federal election data were developed using information from The Canadian Encyclopedia.

Data are quarterly and range from 1978 to 2019. Although data are from quality sources, some issues are still present. The data have limited observations, as monthly data are unavailable for many variables of interest. The data also face limitations because we cannot track individuals' progression over time, primarily their earnings, immigration stream, personal characteristics, employment, etc. Therefore, we cannot attribute unemployment to individual characteristics and instead take a macro approach. To acquire in-depth data, one must work closely with Statistics Canada.

Table 1—Ontario Summary Statistics

Statistic	Min	25th Percentile	Median	75th Percentile	Max
Year	1978	1989	1999	2009	2019
Minimum Wage	8.15	9.16	9.95	11.25	14.27
GDP (Millions $)	976,111	1,174,200	1,470,155	1,950,776	2,317,415
Net Migration	755	15,172	20,745	25,021	41,364
Population	8,604,263	9,969,308	11,452,857	12,998,345	14,637,880
Unemployment Rate	4.83	6.4	7	8.67	12.13
Labour Force	1,412,500	1,799,133	1,991,833	2,360,067	2,648,433

Table 2—Quebec Summary Statistics

Statistic	Min	25th Percentile	Median	75th Percentile	Max
Year	1978	1989	1999	2009	2019
Minimum Wage	8.08	9.16	9.82	11.11	12.93
GDP (Millions $)	930,746	1,130,268	1,451,902	1,970,686	2,411,932
Net Migration	645	4,700	7,322	10,107	14,288
Population	6,442,774	6,882,602	7,315,053	7,843,383	8,540,429
Unemployment Rate	4.97	7.87	9.37	11.37	15.4
Labour Force	957,233	1,136,533	1,224,433	1,402,100	1,539,400

The summary statistics provide a narrative of Ontario and Quebec unemployment rates. From 1978 to 2019, Ontario maintained a higher average and median minimum wage. Ontario also

maintained higher levels of GDP, net migration, population, and labour. However, Ontario has historically displayed trends of lower overall unemployment. Since 1978, Quebec's population has increased by roughly two million, while Ontario's has grown by about six million. In addition to these drastically different growth rates, Ontario has had lower median and mean levels of unemployment.

When analyzing the net migration trends in Canada, we see an overall trend of increasing net migration. The fluctuation in the net migration flow in Canada mimics that of Central Canada, primarily Ontario. The Ontario migration trend varies significantly more than other provinces and territories, exhibiting greater immigration rates than all of Canada (See appendix for summary of trends). As a result, this large influx of immigration in Ontario heavily influences Canada's total migration pattern. Quebec demonstrates the second-highest levels of net migration in all of Canada. British Columbia and Alberta are also showing increasing levels of net migration.

Ontario and Quebec show similar patterns in terms of unemployment rates. However, both Ontario and Quebec demonstrate converging levels of unemployment after 2010. Unlike net migration, Canada's unemployment rate is not driven by unemployment in any province or territory in particular. These trends also do not demonstrate stationarity, which is analyzed in the appendix. Hence for analysis in this paper first differences are taken to ensure data stationarity.

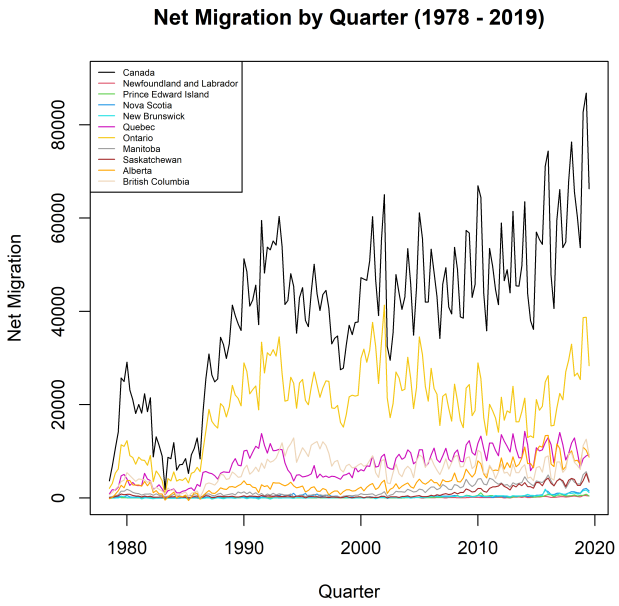

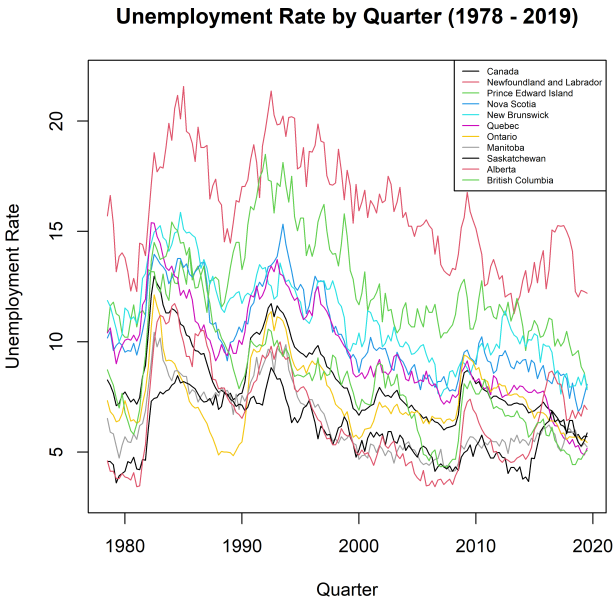

V. Model

To capture the impact of skilled immigration on Ontario's unemployment rate, multiple ARIMAX(p, d, q) models are employed with data for Ontario and Quebec. Although Ontario has the largest net migration levels, Quebec ranks a close second. The primary difference between Ontario and Quebec's immigration policies arise from Ontario's adoption of the PNP system and Quebec's avoidance of it.

An ARIMAX model is utilized as its auto-regressive structure can predict future unemployment rates based on previous rates (importance discussed below). The moving average design of the model can capture shocks in the model. This is important to capture the few periods of high volatility, i.e. recessions. The models take the following form:

$$y_{it} = \beta x_t + \phi_1 y_{t-1} + ... + \phi_p y_{t-p} + \theta_1 \epsilon_{t-1} + ... + \theta_q \epsilon_{t-q} + \epsilon_t$$

where y_t is the covariate of interest, x_t is a covariate at time t and β is its associated coefficient, ϕ_p is the associated coefficient with its autoregressive term y_{t-p} and θ_q is the associated coefficient with its moving average term ϵ_{t-p}.

To capture the variation of the unemployment rate, several variables are incorporated into the model. The ACF and PACF results in the appendix demonstrates persistent autocorrelation and partial autocorrelation. Therefore, the models incorporate lagged unemployment levels. However, the ACF and PACF graphs do not clearly indicate the number of lags. Therefore, simulations determine the optimal number of ARIMAX(p, 1, q) lags; the first difference ensures stationarity. According to Okun's Law, previous levels of unemployment influence future levels

as the unemployment rate eventually settles at a natural level. Therefore, if unemployment is below its natural level, the economy tends to balance itself, exhibiting a rise in unemployment levels, and vice versa.

The labour force should impact unemployment levels directly because the unemployment rate is a ratio between the level of unemployment and the labour force. The unemployment rate declines as individuals leave the workforce due to retirement or injury. The unemployment rate increases as newly graduated individuals enter the workforce.

Net migration is the difference between immigration and emigration numbers. The net migration rate is the ratio between net migration and the provincial population. The net migration covariate and its interaction with other covariates are the interactions of interest. To analyze the interaction between skilled immigration and unemployment rates, the increase in net migration after the introduction of the PNP system is analyzed.

The minimum wage plays a crucial role in determining unemployment levels. As per Okun's Law, unemployment and minimum wages are negatively correlated (as one increases, the other decreases). Throughout the 2000s, Canada saw steady increases in the minimum wage. Therefore, it is essential to incorporate the minimum wage into the model to control for potential employment shocks.

GDP per capita is essential in capturing the state of the economy. If the economy grows more quickly than the population, GDP per capita rises. However, a rapid increase in the population will lead to a decline in GDP per capita. GDP per capita is also negatively correlated with the unemployment rate. As the economy performs better, relative to population growth, unemployment levels decrease.

The party in power is critical for determining immigration levels. The Federal government determines both Canadawide and provincial-level immigration quotas. Liberal governments tend to favour immigration and set higher immigration quotas than their Conservative counterparts. The Trudeau government has continued to show favour towards immigration in recent years and continues to increase immigration quotas yearly.

The unemployment rate graphs demonstrate high levels of volatility during times of recession. Therefore, it is essential to control for these temporary spikes in unemployment with the 1980s, 1990s and 2008 recession covariates.

Indicator variables are incorporated into the ARIMAX model to analyze the impact of Ontario's adoption of the PNP system. These variables include a year indicator variable, denoted by PNP. A second indicator, Ontario, is also utilized to differentiate between Ontario and Quebec. The PNP indicator variable takes a value of 1 if the year is between 2007 and 2019. The Ontario indicator variable takes a value of 1 if the province is Ontario, making Quebec the base province. The interaction between the provincial, year, and net migration rate variables will provide the impact of the skilled immigration policy, PNP, on Ontario's unemployment rate.

VI. Results

Results for ARIMAX(p, d, q) models are shown in table 3. The Akaike Information Criterion (AIC) and Bayesian information criterion (BIC) values suggest utilizing either an ARIMAX(8, 1, 0), ARIMAX(0, 1, 8), ARIMAX(4, 1, 4), ARIMAX(5, 1, 6), or ARIMAX(5, 1, 4).

Table 3—Simulated Models

	Dependent variable:				
	Unemployment Rate$_t$				
	(1)	(2)	(3)	(4)	(5)
Intercept	-0.190^{***}	-0.164^{***}	-0.160^{***}	-0.162^{***}	-0.161^{***}
	(0.043)	(0.029)	(0.033)	(0.032)	(0.032)
Unemployment Rate$_{t-1}$	0.103^{*}		0.237^{**}	-1.119^{***}	0.211^{**}
	(0.053)		(0.101)	(0.091)	(0.088)
Unemployment Rate$_{t-2}$	-0.067		-0.111	-0.247^{**}	-0.062
	(0.044)		(0.106)	(0.099)	(0.100)
Unemployment Rate$_{t-3}$	0.004		0.073	0.070	0.027
	(0.049)		(0.099)	(0.092)	(0.095)
Unemployment Rate$_{t-4}$	-0.295^{***}		0.360^{***}	0.381^{***}	0.410^{***}
	(0.051)		(0.085)	(0.088)	(0.088)
Unemployment Rate$_{t-5}$	-0.042			0.521^{***}	-0.072
	(0.050)			(0.0620)	(0.064)
Unemployment Rate$_{t-6}$	-0.0490				
	(0.043)				
Unemployment Rate$_{t-7}$	0.005				
	(0.046)				
Unemployment Rate$_{t-8}$	-0.220^{***}				
	(0.049)				
ϵ_{t-1}		0.043	-0.198^{***}	1.221^{***}	-0.153^{**}
		(0.054)	(0.075)	(0.104)	(0.066)
ϵ_{t-2}		0.072	0.152^{*}	0.399^{***}	0.109
		(0.059)	(0.080)	(0.084)	(0.070)
ϵ_{t-3}		-0.012	-0.098	-0.034	-0.060
		(0.056)	(0.080)	(0.030)	(0.069)
ϵ_{t-4}		-0.528^{***}	-0.856^{***}	-0.886^{***}	-0.896^{***}
		(0.052)	(0.076)	(0.030)	(0.067)
ϵ_{t-5}		-0.239^{***}		-1.304^{***}	
		(0.053)		(0.098)	
ϵ_{t-6}		0.053		-0.395^{***}	
		(0.056)		(0.089)	
ϵ_{t-7}		-0.050			
		(0.056)			
ϵ_{t-8}		-0.338^{***}			
		(0.050)			
Net Migration Rate	-0.021	0.090	0.067	0.061	0.051
	(0.123)	(0.097)	(0.107)	(0.111)	(0.108)
Minimum Wage	-0.033	0.324^{***}	0.295^{***}	0.293^{***}	0.292^{***}
	(0.112)	(0.053)	(0.062)	(0.068)	(0.064)
GDP per Capita	-0.198^{***}	-0.105^{**}	-0.145^{***}	-0.165^{***}	-0.137^{***}
	(0.069)	(0.047)	(0.052)	(0.052)	(0.053)
Labour Force	0.000	0.000	0.000	0.000	0.000
Federal Party	0.141^{**}	0.092^{***}	0.098^{**}	0.104^{***}	0.097^{**}
	(0.055)	(0.034)	(0.041)	(0.039)	(0.040)
1980s Recession	0.456^{***}	0.427^{***}	0.427^{***}	0.424^{***}	0.423^{***}
	(0.064)	(0.050)	(0.056)	(0.055)	(0.055)
1990s Recession	0.656^{***}	0.566^{***}	0.550^{***}	0.552^{***}	0.553^{***}
	(0.079)	(0.076)	(0.084)	(0.083)	(0.084)
2008 Recession	0.570^{***}	0.659^{***}	0.652^{***}	0.620^{***}	0.663^{***}
	(0.160)	(0.142)	(0.150)	(0.159)	(0.153)
Ontario	0.052	0.011	0.016	0.016	0.015
	(0.052)	(0.024)	(0.031)	(0.029)	(0.029)
PNP	0.158^{***}	0.132^{***}	0.132^{***}	0.142^{***}	0.131^{***}
	(0.057)	(0.034)	(0.040)	(0.038)	(0.039)
Ontario × PNP	-0.046	-0.005	-0.004	-0.002	-0.003
	(0.076)	(0.035)	(0.044)	(0.042)	(0.042)
Ontario × Net Migration Rate	0.201	0.169	0.208	0.213	0.215
	(0.186)	(0.116)	(0.133)	(0.141)	(0.135)
Ontario × Minimum Wage	0.092	-0.205^{**}	-0.196^{*}	-0.197^{*}	-0.191^{*}
	(0.139)	(0.082)	(0.101)	(0.105)	(0.103)
Ontario × Labour Force	0.000	0.000	0.000	0.000	0.000
Ontario × Federal Party	-0.063	0.009	-0.0002	0.001	0.003
	(0.070)	(0.040)	(0.051)	(0.048)	(0.049)
PNP × Net Migration Rate	0.019	-0.098	-0.073	-0.065	-0.056
	(0.125)	(0.100)	(0.110)	(0.113)	(0.110)
PNP × Minimum Wage	-0.143	-0.509^{***}	-0.465^{***}	-0.446^{***}	-0.465^{***}
	(0.190)	(0.118)	(0.134)	(0.132)	(0.132)
PNP × Labour Force	0.000	0.000	0.000	0.000	0.000
PNP × Federal Party	-0.235^{**}	-0.206^{***}	-0.204^{**}	-0.215^{***}	-0.201^{**}
	(0.097)	(0.073)	(0.084)	(0.082)	(0.083)
Ontario × PNP × Net Migration Rate	-0.383^{*}	-0.372^{***}	-0.403^{**}	-0.421^{***}	-0.415^{***}
	(0.208)	(0.141)	(0.161)	(0.162)	(0.160)
Ontario × PNP × Minimum Wage	0.051	0.501^{***}	0.416^{**}	0.330	0.401^{**}
	(0.204)	(0.170)	(0.201)	(0.208)	(0.201)
Ontario × PNP × Labour Force	0.000	0.000	0.000	0.000	0.000
Ontario × PNP × Federal Party	0.112	0.053	0.047	0.047	0.046
	(0.129)	(0.092)	(0.110)	(0.106)	(0.107)
Observations	328	328	328	328	328
Log Likelihood	-120.554	-106.791	-105.339	-100.264	-104.755
σ^2	0.122	0.108	0.108	0.104	0.107
AIC	307.108	279.581	276.678	272.529	277.511
BIC	432.277	404.751	401.848	409.077	412.046

Note: *** p<0.01, ** p<0.05, * p<0.1.

(1) is ARIMAX(8, 1, 0), (2) is ARIMAX(0, 1, 8), (3) is ARIMAX(4, 1, 4), (4) is ARIMAX(5, 1, 6), (5) is ARIMAX(5, 1, 4)

According to the log-likelihood and AIC values, the ARIMAX(5, 1, 6) model has the best fit; however, the BIC values suggest that the ARIMAX(4, 1, 4) model is best. The five models display similar levels of AIC and BIC, indicating support for analyzing all five models. The five models demonstrate similar results with few conflicts.

In all five models, the covariates Ontario and Ontario × PNP return insignificant values. These suggest that before and after the introduction of the PNP system in 2007, both Ontario and Quebec stood at the same level of unemployment, ignoring the impact of other covariates and influences. The PNP covariate suggests that after 2006, both Ontario and Quebec had higher levels of unemployment overall, as all models return significant positive results. These results indicate that after 2006, the unemployment rate increased by roughly 0.131 to 0.158 percent. However, this covariate is likely also capturing some variation from the 2008 recession, given the limited data after 2006. The recession variables also return large significant values for all models.

The models suggest that previous levels of unemployment and unemployment volatility significantly impact future levels of unemployment. All models incorporating autoregressive covariates express that the first lag of the unemployment rate is significant. The direction of influence is inconclusive as we see different signs of the coefficients; three models suggest the impact is positive, and one states it is negative. Models 1, 3, and 5 associate an initial shock of increasing the unemployment rate by one percent with a further increase of 0.103, 0.237 or 0.211 percent in the second quarter. Model 4 suggests the opposite but this result is questionable and the ARIMAX(5, 1, 6), should be analyzed with caution. Overall the impact in a one quarter lag appears to be positive.

The persisting effect of the a shock beyond the second quarter is ambiguous. The models suggest the initial shock's effect at three quarters is negative, although only model 4 returns a significant value. The models suggest the direction of the effect becomes uncertain beyond the second lag as the models do not agree if the impact is positive or negative. The significance of a persisting lag beyond one year (four lags) is also questionable given the lack of significance. Lags of unemployment beyond five quarters seem to show no significance and are likely irrelevant.

The GDP Per capita results are consistent with the theory and the theoretical models. All models return negative coefficients. The smallest prediction suggests that a $1 million CAD increase in GDP per capita will decrease the unemployment rate by 0.105 percent. Intuitively this implies a 0.0000105 decrease in the unemployment rate for every $10,000 CAD increase in GDP per capita. The largest value suggests a 0.198 percent decrease in the unemployment rate for every $1 million dollar CAD increase in GDP per capita.

Regarding the effect of labour on the unemployment rate, all five ARIMAX models return results suggesting insignificance. All coefficients in regard to the labour force are 0, before and after the implementation of the PNP system, in both Ontario and Quebec. The models have determined that other factors explain the variation in the unemployment rate better than the size of the labour force. The labour force trends demonstrate consistent growth as the Canadian population has steadily increased. Alongside this increase in the population and labour force, the Canadian economy has continued to grow to ensure the increasing size of the labour force did not negatively impact the unemployment rate.

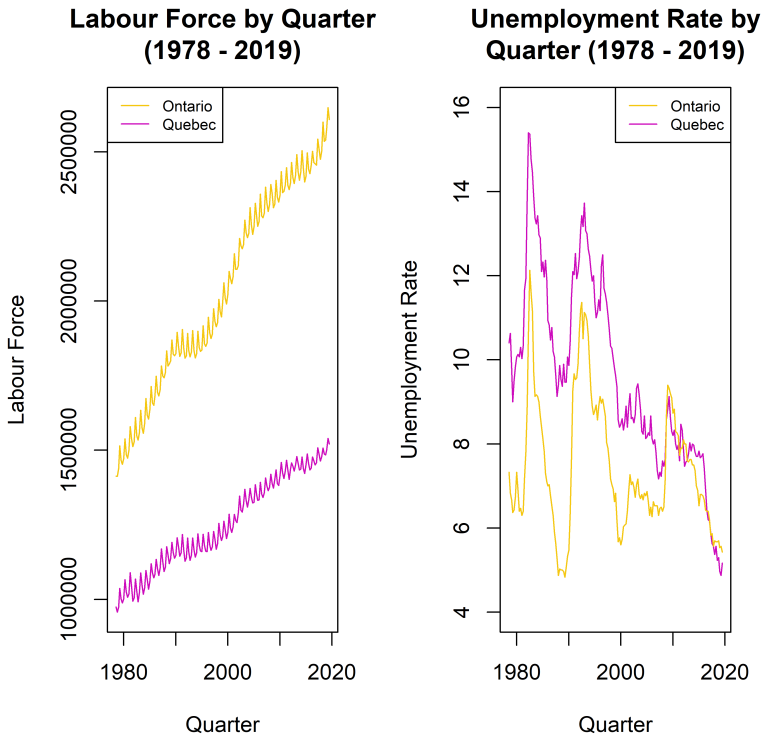

The federal party covariates highlight the impact of the impact of having a Liberal Prime Minister on the unemployment rate. The Canadian Liberal Party tends to favour immigration more and continues to increase immigration quotas. On the provincial level, the results state that the ruling party has no specific impact on Quebec or Ontario individually, but does collectively. The five models imply with significance that before 2007, a ruling Liberal party in the House of Commons would lead to a higher unemployment rate in both Ontario and Quebec. As per the models, this increase ranges from 0.092 percent to 0.141 percent. However, after 2007, this effect has been offset and is now negative. After 2007, a Liberal Prime Minister leads to a decline in unemployment ranging from 0.094 percent to 0.114 percent. However, it is important to acknowledge that the data does not incorporate the period of the COVID-19 pandemic and resulting recession.

Generally the models find that the minimum wage has a significant effect on the unemployment rate. The coefficients suggest that increasing the minimum wage will increase Quebec's unemployment rate more than Ontario's. Before 2007, model 5 states that a dollar increase in Quebec's minimum wage will increase unemployment by 0.292 percent while model 2 suggests a 0.324 percent increase. However, after 2007, the models state that increasing Quebec's minimum wage should decrease unemployment. This result is the opposite of the expected result, as an increase in the minimum wage should always increase unemployment. This result is likely due to the rapid decline in Quebec's unemployment rate after 2009, alongside its rapid minimum

wage increase in the latter decade.

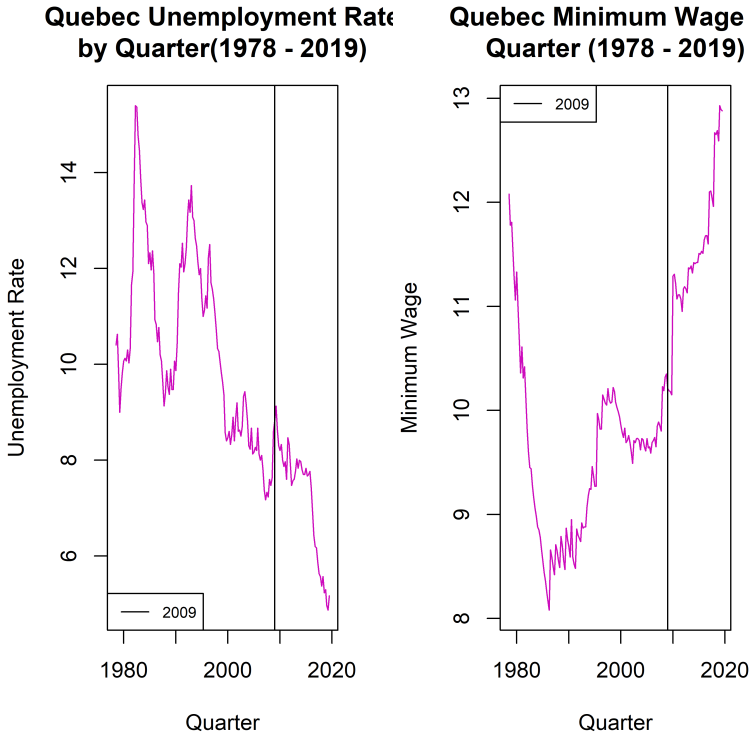

The variable of interest, the net migration rate, captures immigration and emigration. All five models return negative values, suggesting that an increase in the influx of skilled immigration leads to a decline in the unemployment rate. When analyzing its covariates and interactions, all five models return insignificant results except in one situation: Ontario after the introduction of the PNP system. Since the introduction of the PNP system, a one percent increase in the net migration rate will decrease the unemployment rate between 0.372 and 0.421 percent more in Ontario than Quebec. This suggests that after incorporating the PNP system, an increase in net migration leaves Ontario better off than Quebec. These values imply that Ontario's PNP system is more successful at finding skilled foreign labour to improve the provincial economy. The overall consensus seems to be that the PNP system beneficially influences the unemployment rate.

VII. Conclusion

This report analyzed the influence of skilled immigration on unemployment, primarily through the Provincial Nominee Program. By comparing two provinces with similar demographics and immigration systems, this paper aimed to find causation by exploiting differences in provincial

immigration processes. Multiple ARIMAX models were created and selected through simulation. The models unanimously suggest that introducing higher levels of specialized skilled labour based on provincial needs reduces the unemployment rate. This reduction is likely due to the design of the PNP system, which aims to fill labour gaps and initiate new employment streams. However, the model fails to incorporate movement after the end of job contracts. This data limitation makes it difficult to track the exact impact of the PNP program as new immigrants move across Canada. The PNP system aims to find immigrants that will settle in their nominee province, but it does not restrict movement across Canada.

Although the results demonstrate a significant impact, these results are limited to immigration systems using the PNP and points-based system. Therefore, the results may only be applicable in areas that incorporate similar policies to that of Ontario. Other countries that do not utilize strict immigration policies, like Canada, may experience different effects of skilled labour flows, as demonstrated by previous research. The Canadian system aims to use immigration to benefit its economy and selects individuals accordingly. The results suggest that Ontario is successfully incorporating immigration laws to advance the Canadian economy. Selected immigrants do not harm the Canadian economy and provide an overall benefit.

This research can be improved with more precise micro level high frequency data. The models used quarterly data, as not all variables of interest have monthly-level data publicly available. The macro-level data does not allow one to track immigrant movement over time. Therefore as Ontario experiences high levels of migration from other provinces, the model cannot separate Ontario PNP immigrants from Express Entry or other streams. Therefore Ontario not only benefits from skilled immigration inflows from other provinces, but it also intakes more unskilled labour as well. The data does not allow one to analyze the immigration streams individuals follow or their specific skill set and immigration scoring/ranking. Differentiation between the three main immigration streams and their subcategories would help narrow down the impact of immigration streams more accurately. Despite these limitations the results should lend significance to the view that immigration holds net positive impacts on the receiving country or region. This has very topical implications for policy makers.

REFERENCES

[Berry and Soligo 1969] Berry, R A. ; Soligo, Ronald: Some Welfare Aspects of International Migration. In: *Journal of Political Economy* 77 (1969), Nr. 5, S. 778–94. – URL `https://EconPapers.repec.org/RePEc:ucp:jpolec:v:77:y:1969:i:5:p:778-94`

[Borjas 2006] Borjas, G.: Native Internal Migration and the Labor Market Impact of Immigration. In: *The Journal of Human Resources* 41 (2006), Nr. 2, S. 221–258. – URL `https://doi.org/10.3368/jhr.XLI.2.221`

[Canada 2011] Canada, Citizenship: Evaluation of the Provincial Nominee Program. (2011)

[Canada 2022] Canada, Citizenship: Notice Supplementary Information for the 2022-2024 Immigration Levels Plan. (2022)

[Canada 2023] Canada, Citizenship: Canada welcomes historic number of newcomers in 2022. (2023)

[Card 2001] Card, David: Immigrant Inflows, Native Outflows, and the Local Market Impacts of Higher Immigration. In: *Journal of Labor Economics* 19 (2001), Nr. 1, S. 22–64. – URL https://www.jstor.org/stable/10.1086/209979

[Card 2005] Card, David: Is the New Immigration Really so Bad? In: *The Economic Journal* 115 (2005), Nr. 507, S. 300–323. – URL https://doi.org/10.1111/j.1468-0297.2005.01037.x

[Card 2007] Card, David: How Immigration Affects U.S. Cities. In: *CReAM Discussion Paper Series* (2007), Nr. 0711

[Challinor 2011] Challinor, A.: Canada's Immigration Policy: a Focus on Human Capital. (2011)

[El-Assal and Fields 2018] El-Assal, K. ; Fields, D: Canada 2040: No immigration versus more immigration. (2018)

[Gianmarco and Giovanni 2012] Gianmarco, Ottaviano ; Giovanni, Peri: RE-THINKING THE EFFECT OF IMMIGRATION ON WAGES. In: *Journal of the European Economic Association* 10 (2012), Nr. 1, S. 152–197. – URL https://doi.org/10.1111/j.1542-4774.2011.01052.x

[Harris and Todaro 1970] Harris, John ; Todaro, Michael: Migration, Unemployment and Development: A Two-Sector Analysis. In: *The American Economic Review* 60 (1970), Nr. 1, S. 126–142. – URL https://www.jstor.org/stable/1807860

[Islam 2007] Islam, A.: Immigration Unemployment Relationship: The Evidence from Canada. In: *Australian Economic Papers* 46 (2007), Nr. 1, S. 5266. – URL https://doi.org/10.1111/j.1467-8454.2007.00305.x

[Latif 2015] Latif, Ehsan: The Relationship Between Immigration and Unemployment: Panel Data Evidence from Canada. In: *Economic Modelling* 50 (2015), S. 16267. – URL https://doi.org/10.1016/j.econmod.2015.06.013

[Ortega and Peri 2014] Ortega, F. ; Peri, G.: Migration, Trade and Income. In: *Journal of International Economics* 92 (2014), Nr. 2, S. 231–251

[Rohac 2014] Rohac, D.: The Great U.K. Immigration Scare. (2014). – URL https://www.cato.org/publications/ commentary/great-uk-immigration-scare

[Winter-Ebmer and Zweimller 1999] Winter-Ebmer, Rudolf ; Zweimller, Josef: Do immigrants displace young native workers: The Austrian experience. In: *Journal of Population Economics* 12 (1999), Nr. 2, S. 327–340. – URL https://EconPapers.repec.org/RePEc:spr:jopoec:v:12:y:1999:i:2:p:327-340

APPENDIX

A1. Unemployment Rate Stationarity Tests

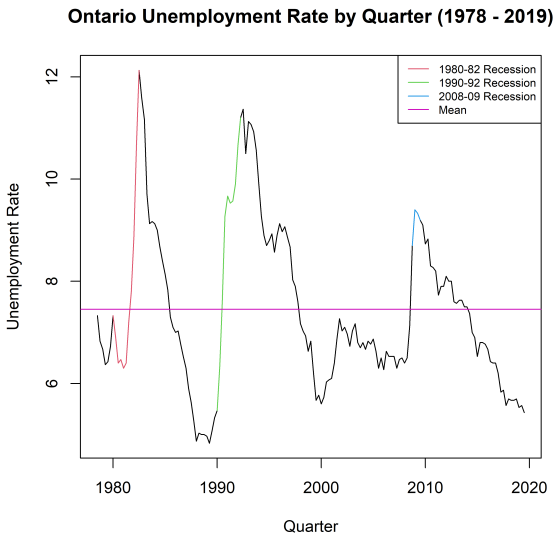

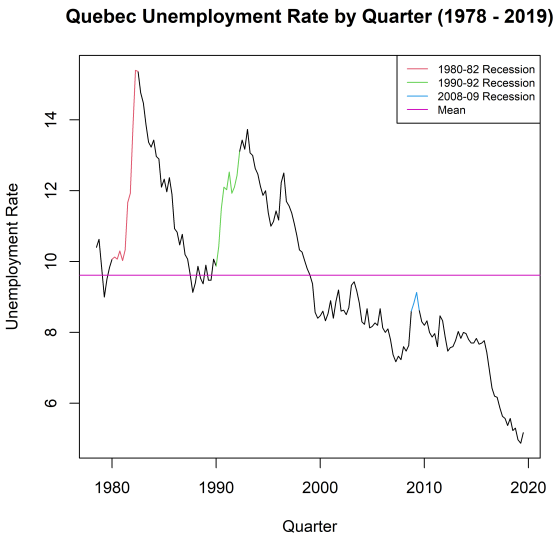

Table A1—Dickey Fuller Test Results

	Ontario		Quebec	
	D.F. Value	P-Value	D.F. Value	P-Value
H_A: Stationary	-1.715	0.695	-2.316	0.444
H_A: Explosive	-1.715	0.305	-2.316	0.555

The Augmented Dickey-Fuller test results indicate that both Ontario and Quebec's unemployment rate trends are non-stationary and not explosive. The figures also clearly indicate the non-stationarity of both series.

A2. First Difference Unemployment Rates

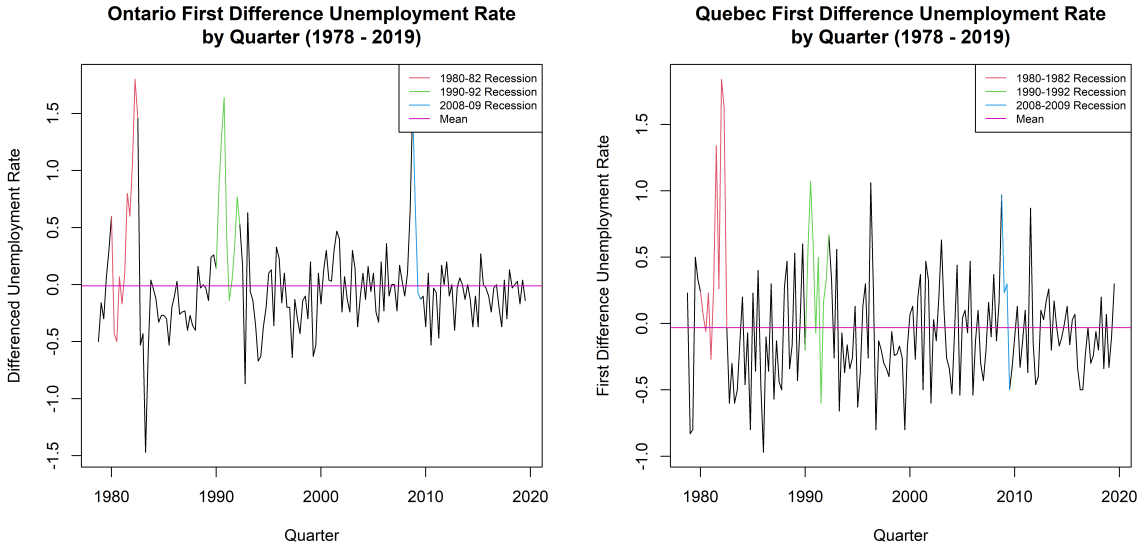

Both Ontario and Quebec's first difference unemployment rate trends are stationary and display occasional spikes during recessions. Due to the limited sample of data, data during these recessions are not dropped and are controlled for. Therefore, MA lags and regression dummy variables are used to control for this volatility.

A3. ACF and PACF Tests for Lag Length

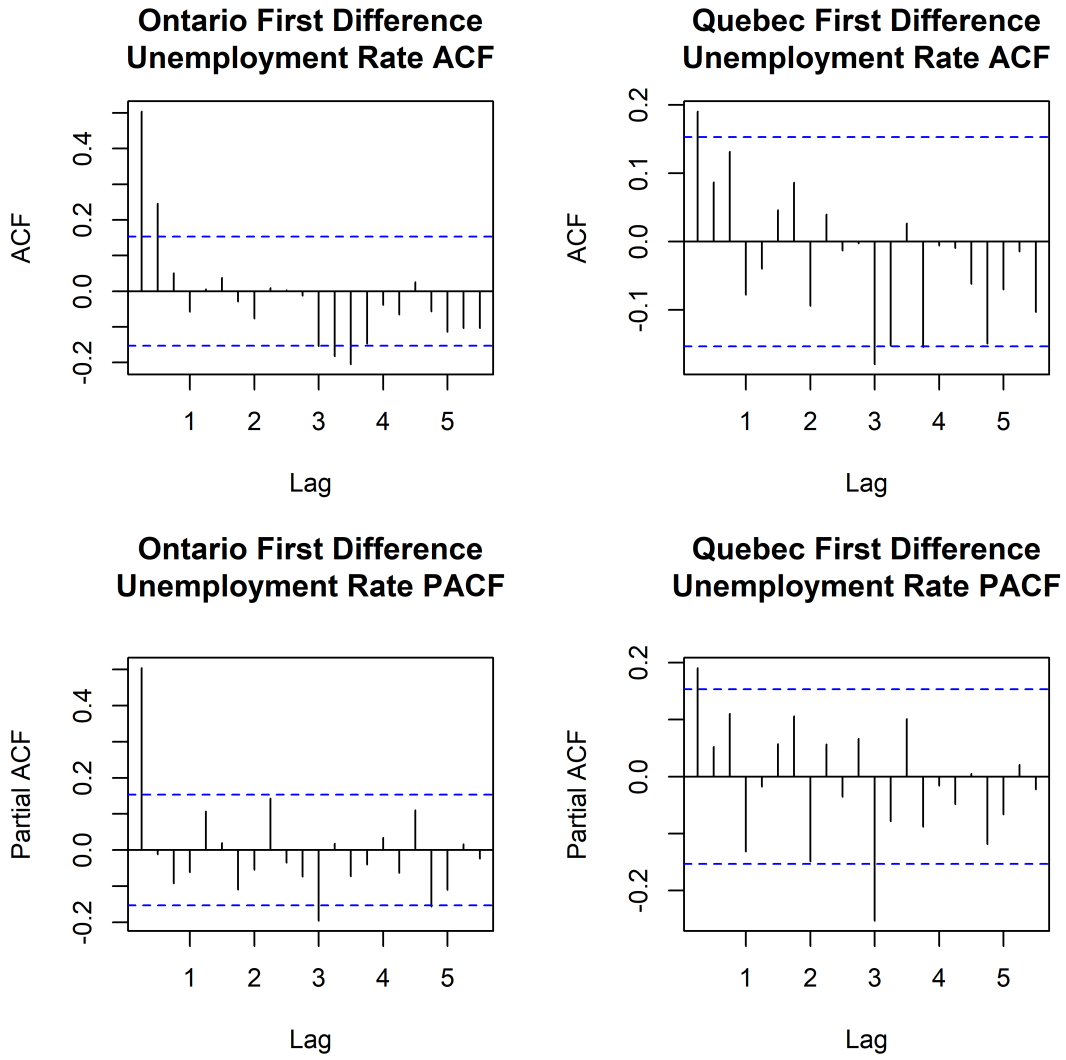

The Ontario ACF persists for three quarter lags but then decays, while its PACF declines after one lag. The Quebec ACF decays after two quarter lags and its PACF declines after one lag. We do not see any clear patterns to suggest the number of AR and MA lags. Therefore, multiple ARIMAX models are designed using simulations to determine the correct number of autoregressive and moving average lags.

A4. ARIMAX Stationarity Tests

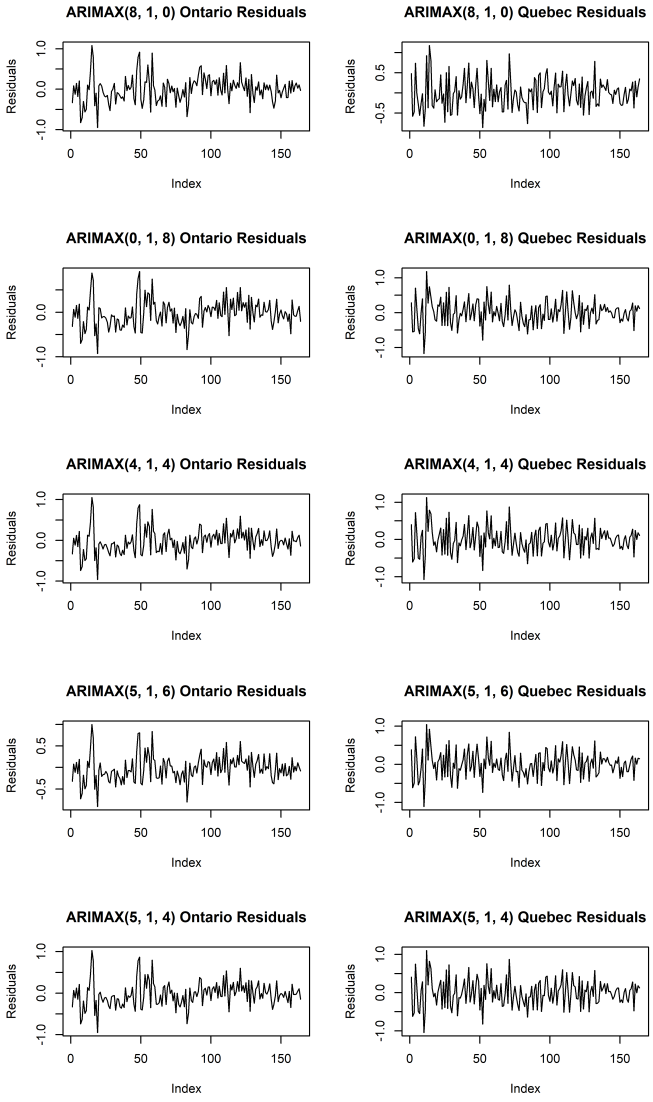

For all models, the residuals appear to be stationary and not explosive. Dickey-Fuller tests were also conducted and confirmed the stationarity of the residuals. The results state that all model residuals are stationary and non-explosive.

Estimating the "Breadwinner Impact" on Household Division of Labour

By Melissa Siqueira*

This study builds upon previous literature examining the determinants of the household bargaining process. By employing a regression-discontinuity design model, I find that breadwinner status can be used as a bargaining tool to avoid household labour. While this effect generally holds for male breadwinners, it does not for female breadwinners. Limiting the sample to just parents, the breadwinner impact differential between males and females becomes negligible, indicating a different dynamic in households with children.

***Keywords**: Regression Discontinuity Design, Household Bargaining, Division of Labour, Family Economics*

I. Introduction

In Canada, the proportion of households with a female breadwinner has been increasing steadily for the last 50 years. Approximately 17.3% of Canadian women earned more than their partners in 2017, compared to just 8% in 1985. Much of this has been attributed to increasing female labour force participation as well as the strong median growth of female wages [Canada (2017)]. Generally, human capital theory states that since time is a limited resource, individuals will choose to divide their time between paid and unpaid work [Becker (1981)]. Moreover, individuals prefer not to do housework and will bargain to avoid engaging in this activity. Given that income is used as a bargaining tool, greater equality in the workplace should be echoed within the home.

Studies have shown that as income increases, housework decreases for male breadwinners. However, this relationship has not held for female breadwinners. All else being equal, women still spend more time than their partners on household labour [Bianchi et al. (2000)]. As such, it is important to determine whether the increasing prevalence of female breadwinners is contributing to the economic advancement of women or if it is putting more burden on these individuals, leading to unsustainable expectations for female-led households. This can have policy implications; if the results support a persistence of unequal division of labour, especially in childcare, it presents a need for a policy that targets this household labour gap.

II. Literature Review

There has been extensive research on division of labour within the household. Becker first modelled division of labour through human capital theory, where individuals divide their finite

* Special thanks to Dr.Nicholas Zammit and Joseph Groga Bada for their invaluable guidance

resource of time between paid and unpaid labour [Becker (1981)]. As individuals prefer leisure to housework, they will bargain to avoid engaging in this activity [Becker (1981)].

The existing literature has attempted to formalize the determinants of this bargaining process. Several studies have looked at the impact of income on division of labour. Bianchi et al. ran an OLS regression model estimating the impact of a wife proportional income on time spent on domestic work, controlling for paid work, employment status, income and education [Bianchi et al. (2000)]. Using time-use survey data from Australia in 1992, they found that as a husband proportion of income increased, so did his wife domestic work. However, they found that the inverse did not hold: in situations where women provided 51-100% of the income, they still engaged in a more traditional division of labour [Bianchi et al. (2000)]. Similarly, Bittman et al. found that income is an important bargaining tool in the division of labour [Bittman et al. (2003)]. Applying an OLS regression model, this time to the American National Survey of Families Households, they predicted a statistically significant negative effect of female income on female housework. Nevertheless this impact was not a one-to-one tradeoff. Thbaud hypothesized that this imperfect correlation between income and housework was due to gender expectations [Thbaud (2010)].

Building upon Thbaud, several studies have attempted to further explore the effect of traditional gender roles on the division of labour within the household [Thbaud (2010)]. Most studies do this by making male to female comparisons, more specifically comparing male and female breadwinners. Chesley et al. analyzed how the breadwinner and caretaker models differ when the roles are reversed between males and females [Chesley and Flood (2003)]. They looked at the American Time Use Survey with cross-sectional data on division of labour between couples, limiting the sample to heterosexual married couples with children where one parent was the breadwinner and their partner a stay-at-home parent. This was their attempt at holding constant all conditions except gender, so that any differences in male and female breadwinner labour should be due to gender effects. Applying an OLS-regression model, they found a statistically significant negative impact of being a breadwinning parent on housework, aligning with previous studies. However, similar to Bianchi et al., they found that division of labour was more equal in female-led households than in male-led households [Bianchi et al. (2000)]. Their conclusion was that males and females in similar economic situations experience divergent outcomes.

Greenstein attempted to quantify gender roles by including a measure of "traditionalism" along with his variable of economic dependence [Greenstein (2000)]. He analyzed panel data from the American National Survey of Families and Households which has information on time use as well as qualitative measures of family attitudes. From this, he created the "traditionalism" dummy which he interpreted as being more entrenched in gender norms. He used seemingly-unrelated regression (SUR) techniques to estimate the impact of a wife's economic dependence (computed as the relative difference in income between an individual and her spouse) and "traditionalism" on household activities. The results show that traditionalism has a positive impact on female hours engaged in housework, and that as economic dependence decreases, household division of labour becomes more equal only until male and female income is approximately equal.

A survey of the literature reveals numerous studies on the effect of income, economic dependence and gender roles on household bargaining via division of labour. However, there have been no recent studies on household division of labour using Canadian data. Moreover, the

majority of the studies have employed standard OLS-regression models. There have been no specific studies on the "breadwinner effect" on the bargaining process and none applying a regression-discontinuity model. Hence this work adds meaningfully to the existing literature by filling several of these gaps.

III. Data and Methodology

A. Data

This paper uses an aggregated version of the General Social Survey on Time Use, a private dataset from Statistics Canada. The Time Use module is a branch of the General Social Survey that focuses on how individuals spend their time working, at home and during social activities. It collects data from individuals aged 15 and over in Canada's 10 provinces via telephone and online questionnaires. Each cycle is a repeated cross-section: individuals are identified by unique codes and do not repeat across periods. All survey questions are linked to actual individual income and household data through tax records. The dataset is an unbalanced pseudo-panel (see Table 1) with the majority of observations coming from the 2015 cycle. For the purpose of this study, the sample was limited to individuals with spouses or common-law partners living in the household, and assumed a dual-earner household in calculation of income ratios. As the goal of this study is to estimate division of labour, it is necessary for an individual to be co-habiting with another individual where their household labour is shared (for example it is not concerned with roommates that live together but lead separate lives). This is in line with the methodology of other division of labour studies.

Table 1—Number of Observations by Year

Year	Number of Observations	Percentage of Total
2005	1,163	6.58%
2010	7,008	39.63%
2015	9,514	53.80%
Total	**17,685**	**100%**

Given each cycle has a different questionnaire, the variables were chosen based on the underlying questions that were unchanged across each cycle. This was done to prevent changes in methodology having an impact on the results. It is also important to mention that the data source is a self-reported survey. As such, the reported duration of activities are proxies of actual duration. It is assumed that the error with which these indicators are measured is random and uncorrelated with any breadwinner status. This is reasonable, as there seems to be no obvious motivation for individuals to under- or over-report their levels of household labour. Some may accidentally report incorrect durations, but this measurement error should not result from any unobserved heterogeneity or because of their breadwinner status. Many other studies have adopted this viewpoint. Reliance on self-reported time use surveys is the standard in division of labour studies. The majority of studies preceding this one have utilized the American

Time Use Survey. The Canadian General Social Survey on Time Use mimics its structure and methodology. This suggests the data will provide consistent results.

The data required various transformations for the purpose of this study. Individual and household income, as well as age, came aggregated by groups rather than displaying exact values. To create the income and age variables, the average value for each group was imputed in real amounts. For simplicity, some background and outcome variables were transformed into binary variables: having a long-term partner living in the household, post-secondary education, gender, immigrant and hiring outside help. Other variables such as number of children and household size were unchanged. Duration of the activites of interest (housework, childcare and paid work) were reported in minutes per day. Finally, the income ratios used to calculate breadwinner status were taken as a proportion of individual income over household income standardized. This was done so that earning half of the income results in a ratio of zero. Further information on derivation of the variables can be found in an appendix available on request.

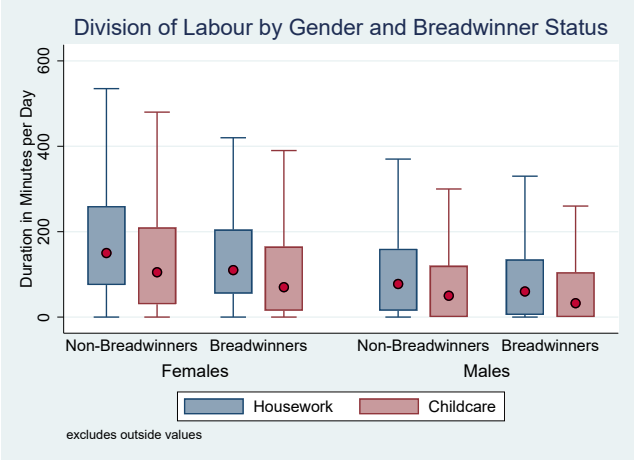

Figure 1. Mean Duration of Household Activities by Breadwinner Status and Gender

Table 2—Total Paid and Unpaid Labour by Breadwinner Status, Full Sample

	Mean	SD	95 % Conf. Int.	N
Non-Breadwinner	357.66	236.52	352.36 - 356.97	7,639
Breadwinner	395.91	256.30	390.92 - 400.89	10,160

As evidenced from a high-level overview of the data, we see that non-breadwinners tend to spend more time on unpaid labour than breadwinners and females more than males (Figure 1). At the same time, breadwinners tend to engage in more total work than non-breadwinners

(Table 2). With this in mind, we might expect to see an inordinately high duration of total labour for female breadwinners.

From the full sample's summary statistics, we can see that there are inherent differences between breadwinners and non-breadwinners that may contribute to the difference in division of labour other than through the breadwinner impact itself (Table 3). Average individual income and post-secondary education are much higher for breadwinners than non-breadwinners, which may indicate that breadwinners have a higher socio-economic status. This could act as bargaining power. There is also a much larger proportion of male breadwinners than female breadwinners. As the literature has found gender to be important in determining household labour, this could explain the labour difference between breadwinners and non-breadwinners.

Table 3—Summary Statistics - Full Sample

	Breadwinners		Non-Breadwinners		Diff in Means
Variable	**Mean (SD)**	**N**	**Mean (SD)**	**N**	**p-value**
Male	0.650 (0.477)	9804	0.269 (0.443)	7649	0.0000
Age	51.2 (14.0)	9582	51.2 (14.9)	7521	0.0148
Post-Secondary Education	0.769 (0.422)	9604	0.661 (0.473)	7981	0.0000
Individual Income	65157 (29162)	9804	24454 (15336)	7639	0.0000
Household Size	2.81 (1.06)	9233	2.82 (1.10)	7089	0.2368
Children	0.55 (0.90)	9804	0.50 (0.88)	7639	0.0001

B. Model

This paper's regression discontinuity model will build upon a popular model within Education Economics, specifically following Lindo et al. study of academic probation [Lindo et al. (2010)]. Their study focuses on estimating the impact of being placed on academic probation on student achievement. They applied a sharp RD model, with their forcing variable being the GPA cutoff and the resulting treatment being placed on academic probation. They then limited their dataset to observations just around the cutoff. The model in this paper takes a similar approach where the income ratio is analogous to distance from a GPA cutoff and the resulting treatment is breadwinner status.

More formally:

$$Y_i = \alpha + \rho D_i + \beta(X_i) + \epsilon_i$$

given:

$$D_i = \begin{cases} 0 & \text{if } incomeratio_i < 0 \\ 1 & \text{if } incomeratio_i \geq 0 \end{cases}$$

where Y_i is time spent on labour (housework, childcare and total work), D_i is a dummy for breadwinner status and X_i is a matrix of covariates including gender, minutes of paid work, household size, whether the individual hires paid help and number of children.

Lindo et al. included their running variable in the regression [Lindo et al. (2010)]. Due to high collinearity (correlation = 0.9692) between the income ratio and breadwinner treatment, including both variables would result in irrelevant variables high with high standard errors. This high correlation arises for multiple reasons. Being a breadwinner is, by definition, a function of the income ratio. However, this is not necessarily an issue as this is a usual setup for an RDD. Given there was no other consistent predictor of becoming a breadwinner, namely because it depends on both individual and household dynamics, only the breadwinner dummy variable was included in the model.

In previous studies, increases in the income ratio have been shown to decrease household labour. Thus, a concern is that excluding the income ratio could result in overestimation of the breadwinner impact. In spite of this, including only the breadwinner status should provide consistent and informative results. Firstly, the sample will be limited to 0.1 units of the income ratio around the cutoff, which results in little variation in the income ratio (only four different values). Essentially, around the boundary we can think of the arising variation not really coming from changes in the income ratio but coming from the change in breadwinner status. Moreover, the goal is to differentiate the impact of the breadwinner impact by gender, which is still possible with this type of analysis. Breadwinner status is included instead of just the income ratio given that the interpretation is easier. It is more difficult to contextualize a movement of one unit of the income ratio than a movement from non-breadwinner to breadwinner status. Results of the models are available with both the income ratio and breadwinner status and just the income ratio in appendix tables A1 and A2. They did not yield results that alter the general conclusions.

The 0.1 unit band on either side of the breadwinner boundary was chosen as it is the smallest boundary that still contained a significant amount of observations for regression analysis. As the original income variables were grouped and discrete, the income ratio is discrete as well. The 0.1 band allows for four different income ratios, two on either side of the boundary. These accounted for approximately 25% of total observations (Table 4).

Table 4—Income Ratio Values, Restricted Sample

Value	Frequency	Percent
-.0714286	1,152	23.93
-.0454545	1,480	30.74
.0555556	893	18.55
.1	1,290	26.79
Total	4,815	100.00

The expanded model interacts the breadwinner variable with gender. Formally:

$$Y_i = \alpha + \rho_1 \text{female breadwinner} + \rho_2 \text{male breadwinner} + \beta(X_i) + \epsilon_i$$

This generates roughly similar estimates to running the model for males and females separately with a breadwinner dummy.

Before running the model, I test whether there is an actual discontinuity in the data. We might see bunching in the breadwinner boundary if individuals negotiate with their employers so they are just earning more than their partner. In this case, being just to the right of the breadwinner boundary would be endogenous; an individual's ambition would cause them to become a breadwinner and this may affect their completion of household work, for example. McCrary recommends a density test of the distribution of the running variable [McCrary (2008)]. As shown in Figure 2, there does not appear to be bunching in the observations immediately to the left or to the right of the breadwinner boundary.

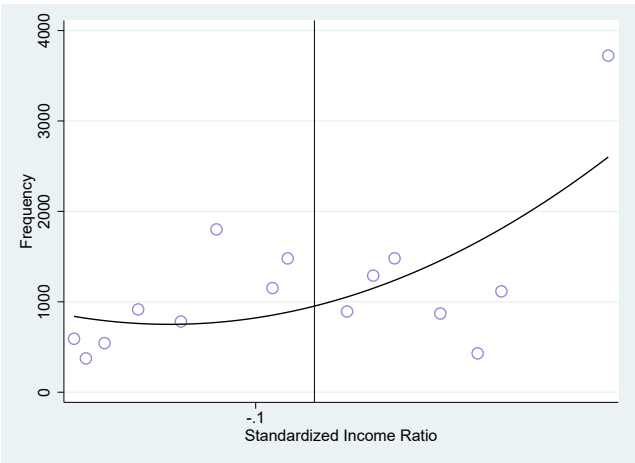

Figure 2. Distribution of Income Ratio

Another check for validity is to check if there is a discontinuity of pre-existing characteristics

at the breadwinner boundary. If the breadwinner boundary is shown to have a statistically significant impact on a pre-determined characteristic, then that leads to questions of validity of the data. As shown in Table 5, the breadwinner boundary is not predicted to have a statistically significant impact on these characteristics.

Table 5—Estimated Discontinuities in Selected Characteristics

Variable	**Age** **(1)**	**Children** **(2)**	**Post-Sec. Educ** **(3)**	**Immigrant** **(4)**
Breadwinner	-0.309	-0.0436	-0.0199	-0.102
	(1.722)	(0.103)	(0.0521)	(0.0723)
Income Ratio	22.50*	-0.209	-0.366	0.896*
	(12.57)	(0.705)	(0.377)	(0.519)
Constant	51.30***	0.513***	0.727***	0.390***
	(0.739)	(0.0443)	(0.0222)	(0.0310)
Observations	4,712	4,815	4,705	2,840
R-Squared	0.009	0.002	0.006	0.001

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

IV. Results

A. Outcomes for All Individuals

I first chose to investigate the impact of the breadwinner status on time spent on housework. As shown in Table 6, the model produces varied estimates of the breadwinner impact on time spent on housework. For all individuals, there is a small impact of a decrease of 8 minutes per day which is significant at the 10% level. When limiting the sample to just males, the breadwinner impact increases to 13 minutes a day and becomes statistically significant at the 5% level. This may seem small, however this decrease represents an hour and a half per week. For females, the model did not predict any statistically significant effect of being a breadwinner with the standard error being larger than the estimated parameter. In the full sample the gender variable has a large statistically significant effect. Being a male is predicted to decrease housework by half an hour a day at the 5% level. As expected, if an individual hires paid help and engages in paid work, they tend to spend less time on housework. Interestingly, having more children isn't predicted to have any significant impact. However, this could be due to a large proportion of the sample not having any children or simply due to the nature of household work changing rather than intensity (note this does not consider changes in childcare but rather other household tasks).

The next outcome investigated was the breadwinner effect on total work in min per day as the sum of housework, childcare and paid work. Preliminary analysis showed that breadwinners, especially female breadwinners, tend to engage in more total labour than non-breadwinners. This may indicate that breadwinner bargaining power is weak when it comes to total work. The regression results do not contradict this high-level view. For all individuals, there is no statistically significant breadwinner impact (Table 7). Similar to the story shown by the regression

Table 6—Regression Results on Housework

Variable	Full Sample (1)	Males (2)	Females (3)	Interaction Effects (4)
Breadwinner	-7.472*	-12.68**	-3.187	
	(4.338)	(6.385)	(5.852)	
Male	-30.45***			-25.91***
	(4.281)			(5.759)
Paid Work	-0.230***	-0.203***	-0.255***	-0.231***
	(0.00726)	(0.0120)	(0.0113)	(0.00820)
Household Size	5.430*	0.713	8.755**	5.262
	(3.208)	(4.928)	(4.318)	(3.250)
Hire	-17.78***	-21.44***	-14.65**	-17.75***
	(4.624)	(7.397)	(6.292)	(4.796)
Children	-4.945	0.505	-9.419*	-4.747
	(3.759)	(6.176)	(5.208)	(3.979)
Individual Income	0.000245	0.000457	6.31e-05	0.000246
	(0.000213)	(0.000321)	(0.000283)	(0.000212)
Female Breadwinner				-2.928
				(5.789)
Male Breadwinner				-13.03**
				(6.397)
Constant	190.7***	159.4***	194.0***	189.5***
	(11.67)	(17.39)	(14.97)	(11.46)
Observations	4,047	1,723	2,324	4,047
R-Squared	0.183	0.153	0.184	0.183

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

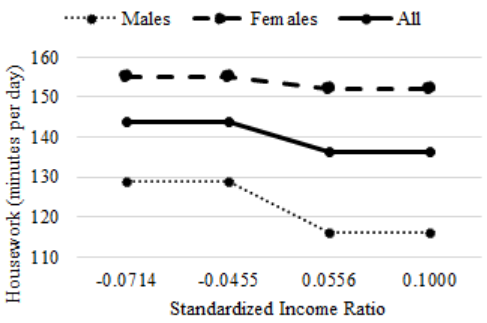

Figure 3. Duration of Housework, All

results on housework, there is a strongly significant impact for male breadwinners, but not for female breadwinners. The predicted male breadwinner impact is a decrease of 25 minutes per day, which is sizeable. Aside from the breadwinner impact, gender is not predicted to have a large or significant impact on total work.

Looking at the other covariates, household size and number of children seem to be key for predicting total work. An increase of one person in the household is predicted to increase total

work by around 30 minutes a day while an increase in one child is associated with a 42 minute increase in total work per day. This makes sense as total work includes childcare which would definitely increase with an extra child. Paid work should also increase as individuals' financial obligations increase.

Table 7—Regression Results on Total Work

Variable	Full Sample (1)	Males (2)	Females (3)	Interaction Effects (4)
Breadwinner	-12.75	-24.09**	-3.974	
	(7.776)	(12.16)	(10.01)	
Male	-19.58**			-9.244
	(7.721)			(10.31)
Household Size	28.77***	22.70**	32.36***	28.37***
	(6.039)	(9.357)	(7.355)	(5.805)
Hire	-13.07	-20.40	-7.046	-12.98
	(8.492)	(14.09)	(10.76)	(8.599)
Children	42.38***	55.14***	34.23***	42.83***
	(7.358)	(11.76)	(8.897)	(7.134)
Individual Income	0.00176***	0.00271***	0.00104***	0.00176***
	(0.000382)	(0.000610)	(0.000483)	(0.000380)
Female Breadwinner				-2.394
				(10.38)
Male Breadwinner				-25.41**
				(11.47)
Constant	240.1***	201.1***	259.1***	237.4***
	(21.07)	(33.12)	(25.57)	(20.55)
Observations	4,047	1,723	2,324	4,047
R-Squared	0.080	0.087	0.073	0.081

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

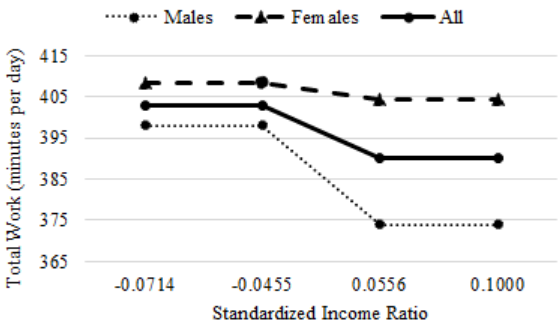

Figure 4. Duration of Total Work, All

B. Outcomes for Parents

After looking at the entire sample around the cutoff, the dataset was then limited to individuals with children (i.e. parents). Three outcome variables were chosen for analysis: housework, childcare and total work. In the previous section, it was decided not to analyze childcare given a large proportion of the sample did not have children; clearly this concern has been addressed by the use of parental data.

Looking at the regression results in Table 8, division of labour seems to be more equal among males and females. There is a similar estimate at the 10% level of the breadwinner impact for both mothers and fathers, a decrease of 20 minutes per day of housework. Note that this impact was stronger than the predicted effect for the general sample. The predicted gender effects are still strongly significant at the 5% level but they are smaller in magnitude with being male resulting in a decrease of housework of only 20 minutes per day compared to over 30 minutes for the full sample. As before, paid work and hiring paid help have important and significant impacts. Household size and number of children remain inconsistently important across the different groups. The R-squared measure also increases in this regression in comparison to the housework regression for all individuals, albeit slightly.

Table 8—Regression Results on Housework, Parents

Variable	Full Sample (1)	Males (2)	Females (3)	Interaction Effects (4)
Breadwinner	-18.13***	-19.05*	-17.82*	
	(6.998)	(10.39)	(9.497)	
Male	-23.03***			-21.97**
	(7.185)			(9.354)
Paid Work	-0.219***	-0.203***	-0.233***	-0.219***
	(0.0135)	(0.0193)	(0.0185)	(0.0134)
Household Size	16.42**	12.25	20.71**	16.37***
	(6.813)	(8.636)	(8.611)	(6.102)
Hire	-20.15***	-14.84	-22.43**	-20.10***
	(6.634)	(10.87)	(9.036)	(6.912)
Children	-15.61*	-16.39	-16.79	-15.57**
	(7.969)	(10.89)	(10.37)	(7.490)
Individual Income	7.34e-05	0.000300	-9.76e-05	7.26e-05
	(0.000353)	(0.000533)	(0.000455)	(0.000346)
Female Breadwinner				-17.08*
				(9.308)
Male Breadwinner				–19.47*
				(10.50)
Constant	172.2***	150.7***	168.5***	172.0***
	(27.16)	(35.32)	(31.47)	(23.39)
Observations	1,236	501	735	1,236
R-Squared	0.225	0.198	0.208	0.225

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

When it comes to childcare (Table 9), the model produces no statistically significant results for the breadwinner effect, which supports the previous literature's claims that childcare is

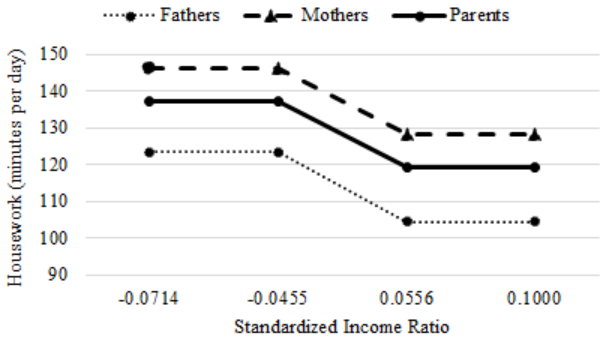

Figure 5. Duration of Housework, Parents

unaffected by the household bargaining process. The main determinants of the amount of time spent on childcare are gender, duration of paid work, household size, hiring paid help and number of children. Here, we see the largest gender difference in division of labour with males predicted to spend on average 35 minutes less a day on childcare. It also is not surprising that a one-person increase in household size is predicted to decrease time spent on childcare as extra individuals within a household may be able to contribute to child-rearing activities such as babysitting.

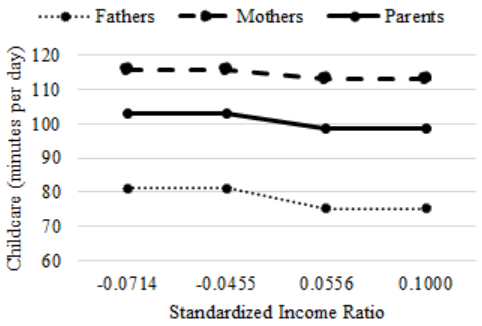

Figure 6. Duration of Childcare, Parents

Finally, the model turns to an analysis of total work for parents. In Table 10, there are hardly any noteworthy results for any determinants in the bargaining process except for a few significant coefficients for household size and children. From the previous regressions, it appears that division of labour is more equal within parenting households so that individuals who engage in more paid labour may have a more proportional decrease in unpaid labour in comparison to non-parenting households. The fact that there are few significant explanatory variables in this regression leads to the R-squared being very low in each of the regression - sometimes less than 1%. Clearly, for parents there are other more important variables that explain total hours worked. Some of these variables could be financial obligations, where they live or future

Table 9—Regression Results on Childcare, Parents

Variable	Full Sample (1)	Males (2)	Females (3)	Interaction Effects (4)
Breadwinner	-4.387	-5.829	-2.769	
	(7.056)	(8.693)	(10.47)	
Male	-35.11***			-31.84***
	(6.694)			(9.466)
Paid Work	-0.105***	-0.0670***	-0.136***	-0.105***
	(0.0134)	(0.0162)	(0.0204)	(0.0135)
Household Size	-25.55***	-21.14***	-29.13***	-25.68***
	(5.133)	(7.225)	(9.493)	(6.176)
Hire	14.43**	24.35***	10.92	14.57**
	(6.884)	(9.090)	(9.962)	(6.995)
Children	44.28***	31.31***	53.24***	44.40***
	(6.292)	(9.110)	(11.43)	(7.581)
Individual income	-0.000289	0.000381	-0.000798	-0.000291
	(0.000355)	(0.000446)	(0.000502)	(0.000350)
Female Breadwinner				-1.130
				(9.420)
Male Breadwinner				-8.546
				(10.63)
Constant	176.1***	102.2***	204.4***	175.5***
	(24.36)	(29.55)	(34.70)	(23.67)
Observations	1,236	501	735	1,236
R-Squared	0.111	0.081	0.089	0.112

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

aspirations that could force them to work more or less.

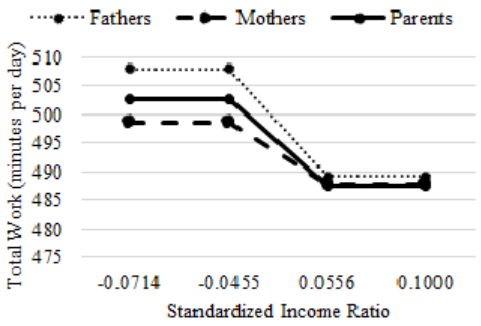

Figure 7. Duration of Total Work, Parents

V. Conclusions

Overall, my results support the existing literature in that income ratio matters - being a breadwinner (having a higher income ratio) is associated in most cases with a decrease in house-

Table 10—Regression Results on Total Work, Parents

Variable	Full Sample (1)	Males (2)	Females (3)	Interaction Effects (4)
Breadwinner	-15.25	-18.82	-10.86	
	(13.95)	(21.83)	(17.57)	
Male	6.321			13.28
	(13.75)			(18.10)
Household Size	-21.78*	-33.05*	-14.30	-22.07*
	(12.23)	(18.10)	(15.93)	(11.92)
Hire	14.63	-9.249	29.73*	14.91
	(13.33)	(22.81)	(16.57)	(13.49)
Children	36.73**	36.67	37.09*	36.99**
	(14.85)	(22.85)	(19.19)	(14.64)
Individual Income	0.000416	0.00207*	-0.000669	0.000411
	(0.000689)	(0.00112)	(0.000842)	(0.000676)
Female Breadwinner				-8.298
				(18.19)
Male Breadwinner				-24.13
				(20.54)
Constant	499.4***	493.3***	506.4***	498.0***
	(46.85)	(72.99)	(57.79)	(45.35)
Observations	1,236	501	735	1,236
R-Squared	0.009	0.016	0.014	0.009

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

hold labour. If household labour was simply determined by the amount of time an individual has in a day, then we would see a one-to-one tradeoff between paid work and household labour. As an individual spends more time on one, then logically they have less time to spend on the other. However, in the majority of the models I ran, this was not the case. Paid work was a statistically significant variable but after controlling for it, there was still a significant impact of the breadwinner impact. This is in line with Becker's human capital theory, which states that individuals use bargaining tools to avoid engaging in unpaid labour (which is an inferior good) [Becker (1981)].

Previous literature has maintained that there are gender inequalities in household division of labour, and the results seem to support this. Behind paid labour, being male had the second-largest statistically significant impact on household labour. These impacts were economically large, ranging from a decrease of 20 to 40 minutes a day. On a weekly basis, this would accumulate to over two hours a week in just gender differences. The impact widened further when interacting the breadwinner impact with gender. Statistically significant impacts of the breadwinner impact on household labour for the full sample became insignificant when accounting for only female breadwinners. The model predicted weak effects on bargaining power for being a female breadwinner.

Another striking feature of the results is how much the results changed when constricting the sample to parents. For housework, the breadwinner impact was statistically significant and of a larger magnitude than for the full sample. Nevertheless, the gender breadwinner differential shrunk as both male and female impacts were equally large and significant. A

possible explanation, hinted at earlier, is that parenting households may tend to be more equal. Preliminary analysis showed that average duration of housework, childcare and paid work tended to be higher for parents. If on average, parents are busier than non-parents, this may leave less room for bargaining power.

A noteworthy result is that paid work had the weakest effect (although still statistically significant) on childcare. In other words, one minute of paid work was worth less in terms of minutes of childcare than in minutes of housework. For childcare, there was essentially no breadwinner impact and unlike the other regressions, household size was calculated to be statistically significant. One of the main facets of Becker's human capital theory is that unpaid labour is undesirable and that individuals will bargain to avoid engaging in this work [Becker (1981)]. However, childcare is of a different nature than housework. It does not necessarily have to be a disagreeable experience, and individuals will often bargain to spend more time with their children. Something else to consider is that childcare cannot necessarily be postponed or delegated like housework. This is why we see household size having such a large impact as additional family members may be trusted to care for children. The lack of significant breadwinner impacts suggest that the bargaining process for childcare differs from that for standard unpaid labour activities.

Although these results are promising there are evidently improvements that could be made to the study. If data were available at the more granular level (i.e. with individual and household incomes instead of income groups), this would allow for a truer RD model with income ratio included in the regression as the running variable. Due to data constraints, the income ratio was discrete and thus had a high correlation with breadwinner status, not allowing the model to separate the two impacts. As well, due to the discrete nature of the data and the lack of availability right at the boundary of the income ratio, it was not possible to completely eliminate omitted variable bias by dropping observations further away from the cutoff. This is evident in how much results changed when including covariates such as paid work and gender.

Another improvement would be to gain access to data with division of labour rather than levels of labour. This survey did not interview individuals within the same household so we do not directly observe division of labour. Some households may naturally have higher levels of labour regardless of their income and education levels and household size. Unfortunately, I was not able to control for this heterogeneity within the regression. Ideally, this study would have been conducted comparing bargaining power between individuals in the same household, as household pairs. This would allow for a more direct study on division of labour.

VI. Implications

Although female breadwinners are on the rise, inequality still persists within households, specifically in division of household labour. In an ideal scenario, individuals would experience a one-to-one tradeoff between their time engaged in paid and unpaid labour. This would allow for a fairer allocation of time - individuals considering working more to increase their income would have to worry less about being overworked as time spent at work would balance out with time spent at home.

Yet, this is not the case. Since breadwinner status has been shown to be a determinant in the bargaining process, an individual may be more inclined to alter their labour force participation

if their income reaches a certain threshold - an amount higher than their partner's. This has implications in game theory for firms offering promotions to managerial roles that offer better pay but require more hours. Not every individual will respond to the same pay raise; it is also dependent on household income and household dynamics.

Another key takeaway is that females tend to take the brunt of unpaid labour, even as breadwinners. As mentioned in the introduction, female breadwinners have been on the rise and there have been many policies aimed at targeting the earnings gap between males and females. Nonetheless, if increasing working hours and household income leads to females increasing their total hours of work, this could lead to female breadwinners being overworked. This has implications for the economy: it may deincentivize females from pursuing breadwinner status hereby hindering female labour force participation.

At a policy level, encouraging individuals to enter the workforce is not simply a question of reducing the wage gap but also attacking the household labour gap. Unpaid labour activities tended to be more inelastic with respect to income for females, especially for childcare. This raises the issue of parents staying out of the workforce due to inability to find or cost of childcare arrangements. Implementing childcare policies or encouraging flex work arrangements within companies could be a way to counteract bargaining differentials within the household.

REFERENCES

[Becker 1981] Becker, Gary S.: *A Treatise on the Family.* Cambridge, MA : Harvard University Press, 1981

[Bianchi et al. 2000] Bianchi, Suzanne M. ; Sayer, Liana C. ; Milkie, Melissa A. ; Robinson, John P.: Is Anyone Doing the Housework? Trends in the Gender Division of Household Labor. In: *Social Forces* 79(1) (2000), S. 191–228

[Bittman et al. 2003] Bittman, Michael ; England, Paula ; Folbre, Nancy ; Sayer, Liana: When Does GenderTrump Money? Bargaining and Time in Household Work. In: *American Journal of Sociology* 109(1) (2003), S. 186–214

[Canada 2017] Canada, Statistics: Household income in Canada: Key results from the 2016 Census / Statistics Canada. 2017. – Report

[Chesley and Flood 2003] Chesley, Noelle ; Flood, Sarah: Signs of Change? At-Home and Breadwinner Parents' Housework and Child-Care Time. In: *Family Relations* 79 (2003), S. 511–534

[Greenstein 2000] Greenstein, Theodore: Economic Dependence, Gender, and the Division of Labor in the Home: A Replication and Extension. In: *Journal of Marriage and Family* 62(2) (2000), S. 322–335

[Lindo et al. 2010] Lindo, Jason M. ; Sanders, Nicholas J. ; Oreopoulos, Philip: Ability, Gender, and Performance Standards: Evidence from Academic Probation. In: *American Economic Journal: Applied Economics* 2(2) (2010), S. 1–26

[McCrary 2008] McCrary, Justin: Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test. In: *Journal of Econometrics* 142(2) (2008), S. 698–714

[Thbaud 2010] Thbaud, Sarah: Masculinity, Bargaining, and Breadwinning: Understanding Mens Housework in the Cultural Context of Paid Work. In: *Gender Society* 24(2) (2010), S. 330–354

APPENDIX

A1. Data Characteristics

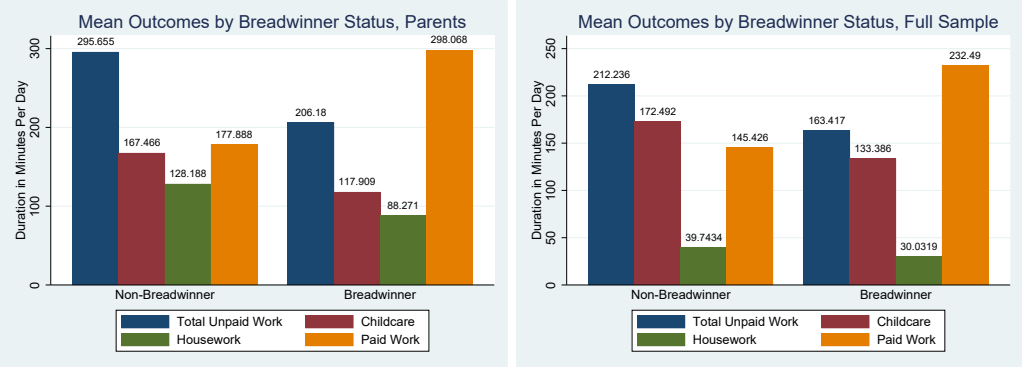

Figure A1. Mean Duration of Labour by Breadwinner Status

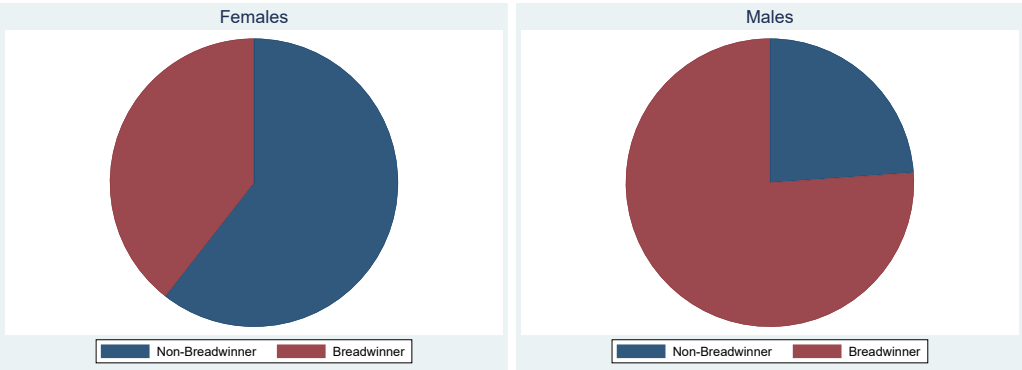

Figure A2. Shares of Breadwinner Status by Gender

A2. Additional Regressions

The base model, without any interaction effects is:

$$Y_i = \alpha + \rho D_i + \beta(X_i) + \gamma incomeratio_i + \epsilon_i \tag{A1}$$

I then expanded the model to capture gender effects:

$$Y_i = \alpha + \rho D_i + \beta(X_i) + \gamma incomeratio_i + \delta incomeratio * gender + \epsilon_i \tag{A2}$$

Results can be found in Table A1. As explained before, the breadwinner variable is directly derived from income ratio, leading to high linear dependence. As such, identification is a problem in this model and this relationship leads to high standard errors and small test statistics. The model did not predict any statistically significant impacts for either variable.

Table A1—Regression Results on Housework Including Breadwinner Treatment and Income Ratio

Variable	**Full Sample** **(1)**	**Males** **(2)**	**Females** **(3)**	**Interaction Effects** **(4)**
Breadwinner	13.08	8.065	15.40	
	(16.62)	(23.60)	(23.23)	
Income Ratio	-155.1	-159.6	-137.8	-153.8
	(119.7)	(169.7)	(167.9)	(119.1)
Male	-30.46***			-25.98***
	(4.282)			(5.759)
Paid Work	-0.230***	-0.202***	-0.255***	-0.230***
	(0.00723)	(0.0106)	(0.00994)	(0.00818)
Household Size	6.060*	1.573	9.113**	5.894*
	(3.208)	(4.592)	(4.383)	(3.235)
Hire	-17.26***	-20.68***	-14.48**	-17.22***
	(4.589)	(6.723)	(6.208)	(4.764)
Children	-5.407	-0.0724	-9.692*	-5.211
	(3.763)	(5.450)	(5.129)	(3.973)
Female Breadwinner				17.40
				(17.19)
Male Breadwinner				7.422
				(17.47)
Constant	190.4***	166.7***	187.9***	189.3***
	(10.84)	(16.09)	(14.58)	(11.03)
Observations	4,047	1,723	2,324	4,047
R-Squared	0.183	0.152	0.184	0.183

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

When the model was run with just the income ratio, excluding the breadwinner impact, it was still possible to find significant results, albeit at the 10% level for the income ratio (recall that there is little variation in the income ratio due to restricted sample size so that standard errors will increase). As shown in Table A2, the most statistically significant and economically large impactful independent variables were gender, duration of paid work and hiring of paid help, as expected. There was no significant impact of the income ratio and gender interaction effect.

Table A2—Regression Results on Housework Using Income Ratio

Variable	**Full Sample** **(1)**	**Males** **(2)**	**Females** **(3)**	**Interaction Effects** **(4)**
Income Ratio	-65.29**	-104.6**	-104.6**	-32.43
	(31.00)	(45.73)	(45.73)	(41.07)
Male	-30.31***			-29.95***
	(4.275)			(4.280)
Paid Work	-0.230***	-0.202***	-0.202***	-0.230***
	(0.00724)	(0.0106)	(0.0106)	(0.00818)
Household Size	5.880*	1.463	1.463	5.701*
	(3.206)	(4.572)	(4.572)	(3.227)
Hire	-17.24***	-20.61***	-20.61***	-17.23***
	(4.587)	(6.719)	(6.719)	(4.764)
Children	-5.257	0.0288	0.0288	-5.052
	(3.762)	(5.432)	(5.432)	(3.968)
Income Ratio ×				-72.53
Gender				(60.94)
Constant	196.1***	170.3***	170.3***	196.9***
	(8.256)	(11.95)	(11.95)	(8.194)
Observations	4,047	1,723	1,723	4,047
R-Squared	0.182	0.152	0.152	0.183

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

The Effects of Unconventional Monetary Policy on Firm Capital Structure

By Jun Takahashi *

This study investigates the effects of unconventional monetary policy on firms' financing decisions using techniques drawn from corporate finance and monetary economics. The data set utilized contains 9,220 firms in 40 countries and 22 sectors from 1998 to 2018. The data is analysed through a variety of techniques including feasible generalized least squares with fixed effects with the Prais-Winsten estimator, reduced vector autoregression, and structural vector autoregression. This research determines that the effects of unconventional monetary policy on capital structure can vary amongst different groups. However, it also finds that most capital structure theories are applicable during periods of unconventional monetary policy. In addition, this work reveals that monetary policy transmission mechanisms differ across conventional and unconventional monetary policy schemes. Unconventional monetary policy most significantly impacts leverage ratios for large private enterprises and has spillover effects globally. The cross-country, cross-sector, cross-firm-type, and cross-firm-size variations suggest that there are group specific factors that determine the impacts of unconventional monetary policy on firm capital structure.

***Keywords**: Capital Structure, Corporate Debt, Conventional and Unconventional Monetary Policy, Global Financial Crisis, Vector Autoregression, Structural Vector Autoregression, Impulse Response Functions*

I. Introduction

After the global financial crisis (GFC) in 2008, major central banks conducted a series of unconventional monetary policy interventions to rescue their economies. In some cases, the policy rate was cut below zero, breaching what was believed to be the lower bound. Conventional monetary policy (CMP) tools reached their limits before fully healing the scar of the global financial crisis. As a result, central banks shifted to unconventional monetary policy (UMP). Although unconventional monetary policies were initially thought to be short-term emergency methods, some remained active for more than a decade. Looking at unconventional monetary policy intervention gives a chance to further explore capital structure decisions. Since monetary policy directly affects the country's interest rate, the policy alters the cost of debt and the opportunity cost of equity. This change in the relative cost of funding will eventually influence the capital structure of firms. Furthermore, the use of unconventional monetary policy can

* I am very thankful to Nicholas Zammit, Xu Tian, Varouj Aivazian, Matthew Liao Mori, Alejandra Zamora Davis, and Jordan Li. Without their extensive support and advice, I wouldn't have been able to complete this research. I also want to thank my friends, colleagues, faculty members, and industry professionals for their insights and feedback.

provide case studies for testing the robustness of existing corporate finance theories. For instance, negative interest rates violate the assumption of the zero-lower bound in the existing literature[1].

This paper is focused on answering three key questions: How does monetary policy affect the capital structure during a crisis and in its aftermath? Is existing capital structure theory still applicable in an extreme state of the world economy? How does the impact of unconventional monetary policy differ from the orthodox monetary policy on corporate financing decisions? Hence, this study interconnects the academic fields of corporate finance and monetary economics. Different techniques from each academic area are applied in this research. In order to address these questions this paper utilizes data from both developed and developing countries with a focus on the effects of UMP that are employed in several high-income countries[2]. The data are organized based on differing monetary policy regimes after 2008, levels of economic development and firm characteristics. Two empirical models are explored, a corporate finance model and a monetary economics model. The framework of the corporate finance model follows [Rajan and Zingales (1995)]; [Booth et al. (2001)]; and [Demirguc-Kunt et al. (2015)]. These studies all utilize a dummy variable for periods of unconventional monetary policy. The monetary economics model employs vector autoregression (VAR) and will be applied to capture the effects of monetary policy and capital structure. The VAR model follows the empirical settings of [Sims (1980)]; [Stock and Watson (2001)]; [Gertler and Karadi (2011)]; and [Lhuissier and Szczerbowicz (2018)].

II. Existing Literature

Past research emphasizes the effects of firm and institutional factors on capital structure. This research suggests how cross-country characteristics such as institutional differences, tax codes, and regulations affect financing decisions [Rajan and Zingales (1995)]; [Booth et al. (2001)]. However, there is limited research on the effects of monetary policy on corporate finance. Similarly, although a large volume of research analyzes the economic impacts of conventional monetary policy, much of the impact of unconventional monetary policy is left to be uncovered.

Since the global financial crisis of 2007/08 firms have been reluctant to rely on long term debt [Demirguc-Kunt et al. (2015)]. Demirguc, Martinez and Tressel found there was a reduction in leverage and debt to maturity in both developed and developing countries from 2008 to 2011. The decline was more severe in the countries with poorer financial infrastructure. These differences resulted from less established banking systems, less sophisticated legal systems around bankruptcy, less investor protection, limited credit information, and more barriers to entry for the banking sector. Large publicly listed firms experienced a smaller decline in their leverage and maturity. The authors concluded that capital markets were a "spare tire" for the large or publicly listed firms since these firms have access to an alternative source of funding when banking systems are damaged.

Classical papers on the international capital structure comparison state that determinants

[1]This paper follows the standard definitions of conventional and unconventional monetary policy from the literature. CMP refers to the modification of key policy rates. UMP includes, but is not limited to, forward guidance (FG), quantitative easing (QE), credit easing (CE), negative interest rates (NIR), and yield curve controls (YCC).

[2]High-income countries are defined by the World Bank as countries with Gross National Income (GNI) per capita higher than US$12,235 in 2016 (calculated based on the Atlas method).

of capital structure are similar among developed and developing countries [Rajan and Zingales (1995)]; [Booth et al. (2001)]. However, there are country-specific factors that influence capital structure. Capital structure is determined by factors such as GDP growth rates and inflation rates which are different for each country. These findings support the view that monetary policy may also be treated as a country specific factor. Although the major central banks set their inflation target at 2%, they must adjust monetary policy depending on country characteristics.

More recent literature on capital structure identifies several key theories including the Static Trade-Off Model (STO), the Pecking-Order Hypothesis (POH), and the Agency Theoretic Framework (ATF). [Myers (1984)] explained the Static Trade-Off as the model that can balance the benefits of interest tax shields and the cost of bankruptcy. The optimal capital structure is found where the marginal benefits of tax shields are equal to the marginal cost of bankruptcy. Using POH, the 'good' companies order the preferred financing source under a financial hierarchy with an internal source (retained earnings or equity), an external debt, and an external equity. The 'bad' companies order the financial hierarchy conversely. The 'good' companies are more profitable than 'bad' companies. If this theory holds during periods of unconventional monetary policy, we shall see an inverse relationship between leverage ratios and the proxy for the firm profitability[3]. The alternative could be explained by the Pecking-Order Hypothesis through the debt-overhang problem [Myers (1977)]. The debt-overhang problem is when companies forego positive net present value projects since the risks are bore by equity holders but benefits are extracted by debt holders. The Agency Theoretic Framework is a model that determines an optimal capital structure given conflicts amongst internal and external investors. The optimal capital structure depends on agency and financing costs. Firms assets and growth opportunities are key determinants for the ATF model [Booth et al. (2001)].

The theoretical consensus on leverage ratios and other factors is summarized in [Harris and Raviv (1991)]. Hariss and Raviv state that "leverage increases with fixed assets, non-debt tax shields, investment opportunities, firm size and decreases with volatility, advertising expenditure, the probability of bankruptcy, profitability and uniqueness of the product." In the presence of asymmetric information, profitability measures are predicted to have a negative relationship with leverage [Myers and Majluf (1984)]. Rajan and Zingales have argued that among the Group of Seven (G7)[4] countries, profitability has a negative impact on leverage [Rajan and Zingales (1995)]. The same is true for developing countries [Booth et al. (2001)]. Booth et al. have shown that the Pecking-Order Hypothesis is appropriate to analyze the various phases of economic development. In other words, the Pecking-Order Hypothesis holds both in developed and developing countries. If POH is applicable among the different states of an economy (i.e. countries with conventional and unconventional monetary policies), then the profitability measure is expected to have an inverse relationship with debt ratios. Fixed assets to total assets (FATA) are assumed to have a positive impact on leverage since the bigger the proportion of tangible assets to total assets, the greater the collateral for the investors. Therefore, the risks of investors bearing agency costs are lower. Higher fixed assets to total assests indicates higher liquidity as well. Hence, FATA is expected to positively correlate with the use of debt.

Literature on the diversity of monetary policy approaches has begun to reveal the relationship

[3]Return on assets (ROA) and gross income are used as proxies for profitability in this research.
[4]G7: Canada, France, Germany, Italy, Japan, the United States, and the United Kingdom

between the choice of policy stance and the capital structure of countries. Several existing unconventional monetary policy models interpret the leverage ratio as one of the key variables during a crisis [Gertler and Karadi (2011)]. Gertler and Karadi have updated dynamic stochastic general equilibrium (DSGE) models of conventional monetary policy with the direct intervention of central banks in capital markets. They successfully quantified the effect of the Federal Reserve System tapping into private markets. Gertler and Karadi's model considers the leverage ratio as an endogenous variable in the dynamism of a steady state where it is profitable for bankers to expand assets in response to an increase in the tolerant leverage ratio of depositors. Their model shows a possible connection between capital structure and monetary policy. Gertler and Karadi also construct a vector autoregressive (VAR) model to study the effect of conventional monetary policy on the economy [Gertler and Karadi (2015)]. They use a monetary policy 'surprise' variable as an instrumental variable for the unobservables in their VAR model. They use the change in interest rates from several monetary policy meeting dates as a policy surprise measure. They find that Forward Guidance (FG) is crucial in a monetary policy transmission mechanism to the real economy. Gertler and Karadi also show that small changes in the short-term interest rates result in large changes in credit costs, which in turn will affect the real economy.

Recently, Gertler and Karadi's framework was repurposed to analyze the effect of unconventional monetary policy [Lhuissier and Szczerbowicz (2018)]. Lhuissier and Szczerbowicz have looked at the effects of monetary policy on aggregate economic activity and firms' debt structure in the United States. They compared the results of VAR models for conventional and unconventional monetary policies. Their findings suggest that conventional and unconventional monetary policies impact capital structure differently. Conventional monetary easing raises the number of loans to non-financial corporations and lowers corporate bonds, whereas unconventional monetary easing is associated with higher bond financing and no change in the number of loans. Conventional monetary policy affects the real economy through interest rates, asset prices, exchange rates, access to credit, and the banking system [Kuttner (2018)]; [Kashyap and Stein (1994)]; [Bernanke and Gertler (1995)]. Unconventional monetary policy influences the real economy through various other channels. Past research identified at least five major transmission channels of unconventional monetary policy.[5] These transmission channels include signalling, portfolio balancing, liquidity, exchange rate, and bank lending [Dell'Ariccia et al. (2018)]. These models and findings will inform the models constructed in this paper.

III. Data Summary

Data has been collected and organized to analyze both the corporate finance and monetary economics models. For the corporate finance model, the firm- and country-level data consist of quarterly data for 9,220 firms from 40 countries over the period 1998 to 2018 (1998 Q4 to 2018 Q4). For the monetary economic model, firm and country level data consists of monthly data for G7 countries over the period 2001 to 2018 (Jan 2001 to Dec 2018). Firm-level data were collected from FactSet. Country-level data were gathered both from FactSet and the World

[5]Notable examples are: [Alsterlind et al. (2015)]; [Amiti and Weinstein (2013)]; [Bernanke (2014)]; [Cahn et al. (2017)]; [Campbell et al. (2012)]; [Curcuru et al. (2018)]; [Fawley and Juvenal (2012)]; [Fiedler et al. (2016)]; [Haldane et al. (2016)]; [Series (2017)]; and [Mokhova and Zinecker (2014)].

Bank data repository.

In both the corporate finance model, three measures of capital structure act as dependent variables in different model specifications. These measures are total debt to total assets (TDTA), long term debt to total assets (LTDTA) and long-term debt to total debt (LTDTD). Independent variables include FATA, return on assets (ROA), sales to total assets (STA), gross income to total assets (GITA), total assets (TA), gross domestic product (GDP) per capita, and dummy variables for years under global financial crisis (GFC) and periods of unconventional monetary policy (UMP). The firm level data are categorized into 6 company types, 3 company sizes[6] and 22 sectors. The country level data are identified based on four income levels[7] and global financial crisis indicators from [Laeven and Valencia (2018)].

In the firm level data the majority of observations represent public firms in developed countries. The sample data includes public companies (89.03%), private companies (5.45%), subsidiaries (4.74%), holding companies (0.57%), non-profit organizations (0.17%) and joint ventures (0.04%). In terms of income classes, 59.37% are high-income, 25.56% are upper-middle-income, 14.32% are lower-middle-income and 0.75% are low-income countries. Small companies make up 11.48%, medium companies 11.35% and large companies 77.17% of all observations. Due to the extreme debt position of some companies, the full sample has been restricted with outliers removed. Both the restricted and unrestricted sample have been indicated on all tables throughout. Table A1 in the appendix provides a breakdown of the data across various categories discussed.

The variables of interest for this research are total debt to total assets (TDTA), long term debt to total assets (LTDTA) and long-term debt to total debt (LTDTD). The first two ratios, TDTA and LTDTA, measure the firms' leverage, which indicates how much of their assets are financed by either short-term debt, long-term debt or both. The third ratio, LTDTD, measures the maturity composition of firm leverage. Total Assets (TA) is the sum of total current assets, long-term notes receivable, total investments and advances, property, plant, and equipment, intangible assets and deferred tax assets. Total Debt (TD) includes short-term debt, the current portion of long-term debt, long-term debt excluding capitalized leases, and capitalized lease obligations. Long-term debt (LTD) consists of any convertible long-term liabilities, bonds, finance leases, long-term royalties, long-term notes payables, preferred liabilities, interest-free loans, borrowing reported as part of total debt, industrial revenue bonds, revolving credit, senior subordinated bonds and notes, and subordinate loans[8].

Table 1) and Table 2) summarize the data under both approaches taken in this paper. The data indicates that considerable dispersion and outliers are visible without restrictions on the sample. Without any restrictions, the global means of TDTA, LTDTA, and LTDTD are 0.3512, 34.41 and 42.1865 respectively. Once TDTA, LTDTA, and LTDTD are strictly restricted to less

[6]The size of the companies is determined by the number of employees aligning with the literature. The small companies have less than or equal to 99 employees. The medium companies have between 100 and 249 employees. The large companies have more than or equal to 250 employees.

[7]Following the World Bank's definitions, country classifications by income levels for 2018-2019 are: low-income if GNI per capita < 996, lower-middle-income if GNI per capita is higher than or equal to 996 and lower than or equal to 3,895, upper-middle-income if GNI per capita is in between 3,896 and 12,055, and high-income if GNI per capita is $> 12{,}055$. GNI per capita is in current USD.

[8]The definitions of TA, TD and LTD differ across industries and countries. The complete lists of constituents in each measure are available from the author on request.

Table 1—Data Summary for the Corporate Finance Model

	Without Restrictions					With Restrictions				
Variables	N	mean	sd	min	max	N	mean	sd	min	max
Firm Characteristrics										
TA	885,119	3,049	21,289	0	$1.529e^{+06}$	482,896	3,913	23,839	0.000115	$1.529e^{+06}$
TD	908,629	716.8	4,446	-2.942	356,571	482,896	953.6	4,990	-2.942	356,571
LTD	568,598	635.7	3,560	-355.7	275,800	482,896	705.1	3,702	-35.02	255,011
ROA	874,903	-12.93	5,482	$-4.622e^{+06}$	292,069	479,764	2.216	182.4	-90,510	15,703
FA	859,895	1,563	9,495	-0.350	546,138	469,632	2,064	11,360	-0.350	546,138
Sales	910,988	1,776	7,998	-143.8	333,542	482,767	2,342	9,615	-0.162	333,542
Cash	908,508	281.2	1,980	-0.455	220,696	482,809	354.8	2,136	-0.455	133,768
Gross Income	576,666	127.4	645.6	-68,336	90,958	455,471	143.2	686.8	-68,336	90,958
TDTA	884,277	0.351	23.50	-0.0128	11,968	482,896	0.254	0.254	-0.0128	9.478
LTDTA	539,094	34.43	17,807	-0.103	$9.253e^{+06}$	482,896	0.134	0.182	-0.103	8.350
LTDTD	497,389	42.20	17,423	-3.235	$8.793e^{+06}$	482,896	0.486	0.467	-3.235	9.997
FATA	837,510	0.565	1.896	-0.00304	807.0	469,632	0.600	2.400	-0.00304	807.0
STA	884,414	1.012	5.641	-0.791	2,386	482,767	0.966	0.805	-0.193	50.18
CTA	884,254	0.168	0.163	-0.325	1.519	482,809	0.154	0.142	-0.325	1.519
GITA	544,756	23.56	12,633	-1,022	$7.993e^{+06}$	455,471	0.0589	1.656	-1,022	227.9
Country Characteristics										
GDP	939,330	25,695	250,869	346.5	$2.177e^{+06}$	253,517	27,299	114,334	373.0	$2.177e^{+06}$
Log GDP	939,330	9.188	1.578	5.848	16.90	253,517	9.620	1.353	5.922	16.90
UMP Dummy	$1.408e^{+06}$	0.168	0.374	0	1	482,896	0.252	0.434	0	1
GFC Dummy	$1.408e^{+06}$	0.0741	0.262	0	1	482,896	0.0591	0.236	0	1
Low-Income	$1.408e^{+06}$	0.00748	0.0861	0	1	482,896	0.00161	0.0401	0	1
Low-Middle-Inc	$1.408e^{+06}$	0.143	0.350	0	1	482,896	0.0989	0.298	0	1
High-Middle-Inc	$1.408e^{+06}$	0.256	0.436	0	1	482,896	0.267	0.443	0	1
High-Income	$1.408e^{+06}$	0.594	0.491	0	1	482,896	0.632	0.482	0	1

Note: Restrictions of TDTA, LTDTA and LTDTD <10.00 are imposed.

Source: FactSet and World Bank Data Repository

than 10.00 then the global means of TDTA, LTDTA and LTDTD are 0.254, 0.134 and 0.485. By excluding the extreme cases of debt usage, analysis can be better focused on the goals of this research. The data summary also indicates a link between income class and other factors. Amongst income classes, companies in upper-middle- and lower-middle-income countries have higher leverage with shorter debt maturities than companies in low- and high-income countries. The private firms have larger leverage with longer deb maturity than the public firms. The larger the companies the smaller the leverage, and the shorter the debt maturity.

Table 1) summarizes the data for the corporate finance model. In this approach there are eight key independent variables to consider. The fixed asset variable represents the company's long-term tangible assets that are mainly used in its operations. Sales for commercial companies includes the sales of goods and services that are discounted by cash, trade costs, sales taxes, and exercise taxes. For financial companies, sales refers to total operating revenue[9]. The return on assets and gross income variables are proxies for firm profitability. The total assets variable

[9]For banks, sales include interest and fees on loans, interest on Federal Funds and bank deposits, lease financing, income from trading accounts, foreign exchange income, investment securities gains/losses, trust income, and commissions. For the insurance companies, sales are premium earned, investment income, and gains/losses on the pre-tax sale of securities.

Table 2—Data Summary for Monetary Economics Model

Variables	N	mean	sd	min	max
country_id	1,944	5	2.583	1	9
year	1,944	2,010	5.189	2,001	2,018
month	1,944	6.500	3.453	1	12
year, month	1,944	599.5	62.37	492	707
one_year	1,664	1.290	1.607	-1.215	6.334
five_year	1,930	2.136	1.661	-1	7.494
GDP	1,944	60,859	162,107	469.0	551,958
CPI	1,943	120.5	14.69	96.23	157.5
mkt_index	1,944	8,092	9,663	48.06	59,715
log CPI	1,943	4.784	0.121	4.567	5.060
Log GDP	1,944	8.322	1.902	6.151	13.22
TDTA	1,944	0.251	0.0663	0.000108	0.995
LTDTA	1,944	0.181	0.0652	0.000108	0.404
LTDTD	1,944	0.684	0.168	0.0307	1.485

Note: Restrictions of TDTA, LTDTA and LTDTD <10.00 are imposed.

Source: FactSet and World Bank Data Repository

reflects the firm size. The GDP per capita variable is used as the indicator of both the economic and institutional development of a country. The dummy variable for a Global Financial Crisis (GFC) takes the value of 1 if the period is between 2008Q2 and 2009Q2, otherwise takes the value of 0. The unconventional monetary policy (UMP) dummy variable assigns 1 for the period after 2008Q2 and assigns 0 otherwise. Each variable's definitions and data sources are available from the author on request. The descriptive statistics obtained from this data are aligned with the existing literature assuring that there are no major issues.

Table 2) summarizes the data for the monetary economics model. In this approach there are several main independent variables. The variables measure inflation, calculated with consumer price index (CPI), nominal GDP, capital structure, interest rates and major financial market indices for each country selected. The capital structure measures are clustered at the country level for each quarter and then interpolated to monthly data. The monthly nominal GDP variable is also interpolated from quarterly nominal GDP. In the VAR model, CPI and nominal GDP are in natural log terms. Following VAR constructions of Gertler and Karadi and Lhuissier and Szczerbowicz, one-year rates and five-year rates of the government bonds are used to calculate the interest rates for both conventional and unconventional periods [Gertler and Karadi (2011)][Lhuissier and Szczerbowicz (2018)]. It is believed that interest rates of the government bonds incorporate the expectations of future interest rates. Due to the fact that unconventional monetary policy targets both short-term rates (one-year rate) and long-term rates (ten-year rate), an average five-year rate is employed as the unconventional policy indicator. Financial market index for each country is: S&P/TSX for Canada, CAC40 for France, DAX for Germany, FISE MIB for Italy, Nikkei225 for Japan, FTSE100 for the UK and S&P500 for the US. Table 2 shows the summary statistics for the monetary economics variables.

IV. Model Design

The corporate finance and monetary economics approaches each examine economic variables differently. While the corporate finance approach disentangles the cross-sectional variations in unconventional monetary policy on capital structure, the monetary economics approach reveals the structural differences in dynamic effects of conventional and unconventional monetary shocks on corporate financing decisions.

For the corporate finance approach, a simple linear regression with various capital structure measures, firm characteristics and country characteristic data are constructed. Since capital structure is persistent, the presence of the autocorrelation across periods makes ordinary least squares estimations biased. Instead, a feasible generalized least squares estimation with fixed effects and Prais-Winsten estimators for the serially correlated error terms is employed. Two time dummy variables are included in the regression model to account for the effects of both the global financial crisis and the introduction of unconventional monetary policy. The simple regression model is the following:

$$(1)\quad Y_{ijt} = \alpha + \beta \cdot FirmControls_{ijt} + \gamma \cdot CountryControl_{jt} + \mu_0 \cdot GFC_t + \mu_1 \cdot UMP_t + f_i + \epsilon_{ijt}$$

where Y_{ijt} is a capital structure measure, either TDTA, LTDTA or LTDTD for firm i in country j at time t. $FirmControls_{ijt}$ are FATA, ROA, STA, GITA, and TA. $CountryControl_{jt}$ is the natural log term of GDP per capita. GFC_t and UMP_t are dummy variables. f_i is the fixed effect within each sample of firms and countries. ϵ_{ijt} is the error term. The error term is white noise and assumed to follow a first order autocorrelation. The coefficients of interest are μ_0 and μ_1.

For the monetary economics approach, impulse response functions are analyzed using a simple reduced VAR and structural VAR (SVAR). The model includes variables measuring the CPI, GDP, capital structure, interest rates, and financial market indices for the G7 countries. The country selections are based on the operation of their monetary policy. The reduced-form VAR model is the following:

$$(2)\quad Y_t = \sum_{i=1}^{\rho} B_i Y_{t-i} + \alpha_y + \epsilon_t,\ t = 1, 2, ..., T,$$

where Y_t is an $n \times 1$ vector of endogenous variables at time t, B_i is an $n \times n$ coefficient matrix, Y_{t-i} is the i lagged variable of Y_t, α_y is the constant term of an endogenous variable y, ρ is the number of lags, ϵ_t is the forecast error term[10], n is the number of endogenous variable, and T is the sample size. The general form of the SVAR is given by:

[10]The forecast error term, ϵ_t, is assumed to be white noise

$$AY_t = \sum_{i=1}^{\rho} C_i Y_{t-i} + \epsilon_t,\ t = 1, 2, ..., T \tag{3}$$

where A is an $n \times n$ coefficient matrix capturing the contemporaneous relationships among variables in Y_t, C_i is an $n \times n$ coefficient matrix, and the rest matches equation 2. For the post estimation analysis, the error term is transformed to a linear combination of mutually orthogonal 'structural' shocks. Assuming the variance-covariance of the error terms to be an identity matrix, let B be the $n \times n$ identifiable matrix:

$$\epsilon_t = Bu_t \tag{4}$$

$$E\big[u_t u_t'\big] = I. \tag{5}$$

By substituting the equation 4 into the equation 3, the SVAR becomes:

$$AY_t = \sum_{i=1}^{\rho} C_i Y_{t-i} + Bu_t,\ t = 1, 2, ..., T \tag{6}$$

where u_t is the linearly independent structural shocks. Assuming A is invertible, the reduced-form of SVAR can be obtained by multiplying the equation 6 by A^{-1}.

$$Y_t = \sum_{i=1}^{\rho} A^{-1} C_i Y_{t-i} + A^{-1} B u_t \tag{7}$$

Following 'recursive' identification, A is set to be a unit lower-triangular matrix and B is set to be a diagonal matrix. Cholesky ordering sets the endogenous variables in Y_t based on the impact that the endogenous variable has on other endogenous variables. The most influential endogenous variable comes first and the least influential endogenous variable comes last. In the case of this research, the endogenous variables are ordered as CPI, GDP, debt measures, interest rates, and then financial market indicators. The impulse response function can be derived from equation 7 using the lag operator L (which is defined as $L^p x_t = x_{t-p}$). The impulse response function for SVAR is:

$$Y_t = (I - A^{-1} \sum_{i=1}^{\rho} C_i L^i)^{-1} A^{-1} B u_t. \tag{8}$$

Since the focus of this research is the effects of unconventional monetary policy on capital

structure, VAR models are run on two different periods for conventional and unconventional monetary policy. Each country has different conventional and unconventional periods. Using equation 6 and assumptions about matrices A and B, the SVAR model in this research is constructed as follows:

$$(9)\quad \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ -a^{21} & 1 & 0 & 0 & 0 \\ -a^{31} & -a^{32} & 1 & 0 & 0 \\ -a^{41} & -a^{42} & -a^{43} & 1 & 0 \\ -a^{51} & -a^{52} & -a^{53} & -a^{54} & 1 \end{bmatrix} \begin{bmatrix} CPI_t \\ nGDP_t \\ \textit{Capital Structure Measure}_t \\ \textit{Interest Rate}_t \\ \textit{Financial Market Index}_t \end{bmatrix} =$$

$$\sum_{i=1}^{\rho} \begin{bmatrix} c_{t-i}^{11} & c_{t-i}^{12} & c_{t-i}^{13} & c_{t-i}^{14} & c_{t-i}^{15} \\ c_{t-i}^{21} & c_{t-i}^{22} & c_{t-i}^{23} & c_{t-i}^{24} & c_{t-i}^{25} \\ c_{t-i}^{31} & c_{t-i}^{32} & c_{t-i}^{33} & c_{t-i}^{34} & c_{t-i}^{35} \\ c_{t-i}^{41} & c_{t-i}^{42} & c_{t-i}^{43} & c_{t-i}^{44} & c_{t-i}^{45} \\ c_{t-i}^{51} & c_{t-i}^{52} & c_{t-i}^{53} & c_{t-i}^{54} & c_{t-i}^{55} \end{bmatrix} \begin{bmatrix} CPI_{t-i} \\ nGDP_{t-i} \\ \textit{Capital Structure Measure}_{t-i} \\ \textit{Interest Rate}_{t-i} \\ \textit{Financial Market Index}_{t-i} \end{bmatrix} + \begin{bmatrix} b_t^{11} & 0 & 0 & 0 & 0 \\ 0 & b_t^{22} & 0 & 0 & 0 \\ 0 & 0 & b_t^{33} & 0 & 0 \\ 0 & 0 & 0 & b_t^{44} & 0 \\ 0 & 0 & 0 & 0 & b_t^{55} \end{bmatrix} \begin{bmatrix} u_t^1 \\ u_t^2 \\ u_t^3 \\ u_t^4 \\ u_t^5 \end{bmatrix},$$

where a^{zw} is the estimated coefficients to account for the contemporaneous relationships between the variable z and the dependent variable in equation w, c_{t-i}^{zw} is the estimated coefficients for the i lagged variable z in equation w, ρ is the number of lags, and b_t^{zw} is the coefficient for the forecast term u_t^z. For z and w, the number 1, 2, 3, 4 and 5, refer to CPI, nominal GDP, capital structure measure, interest rate, and financial market index, respectively. The capital structure measure is either TDTA, LTDTA or LTDTD. The interest rate is the one-year rate and five-year rate of government bonds for the conventional and unconventional periods, respectively. The financial market index varies amongst countries. The lag selections are based on Akaike's[11] and Bayesian[12] information criteria. The number of lags used in SVAR are different across monetary policy schemes and countries. The number of lags used for each SVAR are given in the appendix.

V. Results

The corporate finance model and monetary economic model both investigate the effects of unconventional monetary policy on how firms are financed from different approaches. The variation

[11] According to [Akaike (1998)], Akaike's information criteria is defined as

$$AIC = -2\ln L + 2K$$

where L is the maximized log-likelihood of the model and K is the number of estimated parameters.

[12] According to [Schwarz (1978)], Baysian information criteria is defined as

$$BIC = -2\ln L + K\ln N$$

where L is the maximized log-likelihood of the model, K is the number of estimated parameters, and N is the sample size.

in unconventional monetary policy effects from corporate finance method sheds light on potential explanations relating monetary policy effects to capital structure during a global financial crisis and in its aftermath. Cross-country variations assess the portability of capital structure theories after the intervention of unconventional monetary policy. The impulse response functions from the monetary economic method provides insight into the structural difference between the conventional and unconventional shocks on corporate financing decisions. Each will be discussed in turn.

A. The Effects of Monetary Policy During Crisis

The results of the estimated effects of both the global financial crisis and unconventional monetary policy at the global level, using the corporate finance approach, are presented in table 3. The estimated effects of the global financial crisis (GFC) and unconventional monetary policy (UMP) on capital structure for different groups are presented in columns 9 and 11 of table 3.

Table 3 indicates that both the crisis and unconventional policy affected the firms' capital structure. While the global financial crisis induced firms to deleverage their long-term capital, UMP encouraged firms to take more debt in the short and long term. These results conform to the existing literature. When uncertainty rose significantly after the global financial crisis, many firms were reluctant to take more long-term debt. What is profound from table 3 is that TDTA increased during the crisis periods and the effects of unconventional monetary policies are larger than the effects of the crisis on capital structure in general. The decrease in LTDTA and LTDTD are 0.2 basis points (bps) and 1.65 bps, respectively. TDTA increased by 0.163 bps during the crisis. Even after a decade following the global financial crisis and the introduction of unconventional policy, the results suggest that unconventional monetary policy increased corporate leverage above pre-crisis levels. This finding might partially explain why some major central banks exited or are exiting from unconventional policy. The overuse of unconventional monetary policy might lead to excessive risk-taking behaviour by firms. Mitigating these risks by exiting from unconventional policy seems a reasonable monetary policy response.

The GFC impacted both TDTA and LTDTA for countries that faced a systemic banking crisis by 0.507 bps and 0.805 bps. There was a spillover effect of a decrease in long-term debt to the countries that did not experience a systemic crisis. On average, firms lowered their LTDTA by 0.534 bps and LTDTD by 2.36 bps even for countries that did not suffer a banking crisis. This spillover effect demonstrates the interconnectivity of the banking system and activities around the world. Unconventional monetary policy, on the other hand, had a larger impact on the countries that did not face a banking crisis. The estimated unconventional policy effect on TDTA, LTDTA, and LTDTD for the non-crisis countries was 2.94 bps, 1.29 bps, and 1.67 bps. For the crisis countries these effects were 0.991 bps, 0.167 bps, and 0.208 bps respectively. Unconventional monetary policy had a statistically significant effect on all leverage measures for non-crisis countries. This result could indicate that additional capital went to non-crisis countries as the return on capital was higher in these countries than in crisis countries. Most of the countries that had a banking crisis went into recession and interest rates were artificially low to stimulate the economy. The lower interest rates should have encouraged firms to borrow more, however, the rapid increase in uncertainty in the economy discouraged the global financial crisis countries to take on more debt.

Table 3—Estimation Results on Unconventional Monetary Policy for the Global Financial Crisis

Variable	Obs	Mean	Std Dev	Min	Max	Obs	GFC		UMP		FE
							Est	Std Err	Est	Std Err	
Global											
Global Average											
td/ta	884,594	0.351253	23.49294	-0.01282	11967.59	270,270	0.000866	(0.00414)	-0.0152**	(0.00751)	Yes
ltd/ta	539,225	34.41963	17804.39	-0.1035	9252552	242,463	0.00800	(0.00505)	0.00732*	(0.00380)	Yes
ltd/td	497,493	42.1865	17421.59	-3.23467	8792784	218,125	30.68	(65.55)	-48.47	(54.51)	Yes
Global Average with Restrictions											
td/ta	482,996	0.254	0.254	-0.0128	9.478	217,666	0.00163**	(0.000797)	0.0252***	(0.00157)	Yes
ltd/ta	482,996	0.134	0.182	-0.103	8.350	217,666	-0.00221**	(0.000892)	0.00958***	(0.00137)	Yes
ltd/td	482,996	0.486	0.467	-3.235	9.997	217,666	-0.0165***	(0.00393)	0.0103***	(0.00329)	Yes
GFC Experience											
without GFC shock											
td/ta	372,378	0.241	0.206	-0.0128	8.835	155,867	0.000209	(0.000753)	0.0294***	(0.00153)	Yes
ltd/ta	372,378	0.106	0.141	-0.103	8.350	155,867	-0.00534***	(0.000863)	0.0129***	(0.00151)	Yes
ltd/td	372,378	0.398	0.420	-3.235	9.960	155,867	-0.0236***	(0.00418)	0.0167***	(0.00384)	Yes
with GFC shock											
td/ta	110,518	0.296	0.371	5.24e-07	9.478	61,799	0.00507**	(0.00215)	0.00991**	(0.00390)	Yes
ltd/ta	110,518	0.230	0.256	0	8.316	61,799	0.00805***	(0.00236)	0.00167	(0.00304)	Yes
ltd/td	110,518	0.782	0.496	0	9.997	61,799	-0.00261	(0.00925)	0.00208	(0.00669)	Yes
Income Classes											
Low											
td/ta	779	0.171	0.147	1.51e-05	0.669	556	-0.00474	(0.00817)	0.00561	(0.0161)	Yes
ltd/ta	779	0.128	0.133	0	0.563	556	-0.0132	(0.0106)	0.0117	(0.0181)	Yes
ltd/td	779	0.636	0.396	0	4.903	556	0.120	(0.0741)	0.0719	(0.0565)	Yes
Lower Middle											
td/ta	47,742	0.276	0.237	-0.0128	5.915	15,582	0.000301	(0.00388)	0.0220***	(0.00651)	Yes
ltd/ta	47,742	0.135	0.188	-0.103	5.606	15,582	0.0139**	(0.00665)	-0.0120	(0.0114)	Yes
ltd/td	47,742	0.454	0.461	-3.235	9.917	15,582	-0.00959	(0.0285)	0.00509	(0.0409)	Yes
Upper Middle											
td/ta	129,088	0.256	0.224	4.82e-06	8.835	56,730	0.00238	(0.00159)	0.0218***	(0.00292)	Yes
ltd/ta	129,088	0.0835	0.126	0	8.350	56,730	-0.00849***	(0.00132)	0.0219***	(0.00194)	Yes
ltd/td	129,088	0.298	0.403	0	9.953	56,730	-0.0183***	(0.00662)	0.0250***	(0.00617)	Yes
High											
td/ta	305,387	0.250	0.269	-0.0120	9.478	144,798	0.00248***	(0.000956)	0.0194***	(0.00193)	Yes
ltd/ta	305,387	0.156	0.196	0	8.316	144,798	0.00145	(0.00112)	0.00656***	(0.00160)	Yes
ltd/td	305,387	0.570	0.470	0	9.997	144,798	-0.0159***	(0.00469)	0.0139***	(0.00366)	Yes

Note: *** p<0.01, ** p<0.05, * p<0.1.

Source: FactSet and World Bank Data Repository

Unconventional monetary policy was statistically significant for TDTA in lower-middle income countries and for all capital structure measures in the upper-middle and high-income countries. Firms in upper-middle income countries lowered their long-term debt during the crisis but raised their TDTA, LTDTA, and LTDTD during unconventional policy periods. As the unconventional monetary policy effect on these firms' leverage ratios was larger than the deleveraging effect on long-term debt from the crisis, the firms in the upper middle-income countries were taking on greater leverage than their pre-crisis levels. When the default probabilities were rising for the companies in the upper-middle income countries, investors might have taken the capital out from these companies or banks hesitated to lend capital to these companies during the crisis. However, when the global economy started recovering from the global financial crisis, investors who sought higher yields than the low-interest rates available in the high-income countries might

have invested more into those riskier countries.

For the case of high-income countries, both the crisis and unconventional policy had positive effects on capital structure measures. Unconventional monetary policy influenced the firms that were located in the high-income countries to increase their leverage in the short and long term by 1.94 bps and 0.656 bps. Their debt maturity increased by 1.39 bps as well. It is surprising to see that unconventional policy effects are bigger on TDTA than LTDTA. Following the theory of 'filling-gaps', unconventional policy should have a larger impact on LTDTA than TDTA. The 'filling-gaps' hypothesis argues that while the central bank is acting as a long-term liquidity provider by purchasing long-term debt securities, firms issue more long-term debt to meet the higher demand. Further research is required to understand this contradiction between the theory and the empirical findings.

There are variations in the effects of the global financial crisis and unconventional monetary policy across countries. Some countries endured both, either or neither of the crisis or unconventional policy effects. Both signs and the direction of the effects for the crisis and unconventional policy vary by country. The countries that had statistically significant unconventional monetary policy effects on all leverage ratios are Brazil, China, France, Japan, and the Philippines. Countries with significant unconventional monetary policy effects on both TDTA and LTDTA are Canada, Chile, Czech Republic, Estonia, Mexico, and Turkey. The countries with statistically significant unconventional monetary policy effects only on TDTA or LTDTA are Argentina, Bangladesh, Finland, Germany, India, Ireland, Netherlands, New Zealand, Portugal, South Korea, and Spain. As studies on the effect of unconventional monetary policy on the real economy and firms' capital structure are still in the developmental stage, the factors behind these cross-country variations are ambiguous. However, the presence of these variations shed light on the need for further investigation on why different countries experienced different impacts from unconventional policy. These findings might contribute to the literature by explaining why countries follow different recovery pathways from financial crises. [Demirguc-Kunt et al. (2015)] have researched the cross-country factors for the global crisis effects on firms' capital structure. They found that the crisis effects are more significant in countries with less sophisticated financial infrastructure, weaker banking systems, less sophisticated legal systems, fewer protections for investors, limited availability of credit information and higher barriers for bank entries.

Unconventional monetary policy affected private, public and holding companies but not other types of companies. All capital structure measures for the public companies are affected by unconventional policy, which increased TDTA, LTDTA, and LTDTD by 2.52 bps, 1.02 bps, and 1.12 bps respectively. For holding and private companies, unconventional policy encouraged these companies to take on more debt for the short and long term. In fact, TDTA and LTDTA rose by 5.19 bps and 4.98 bps for holding companies and by 4.37 bps and 4.49 bps for private companies. Unconventional monetary policy had a greater impact on private companies' capital structure for both TDTA and LTDTA than the impact on public companies' capital structure. Considering this result, the higher debt of private companies could be related to the fact that private companies experienced larger reductions in these capital structure measures during the global financial crisis. In order to recover their leverage ratios to the pre-crisis levels, private companies needed to increase their leverage during unconventional policy periods more than

public companies. Looking at this phenomenon through the policy perspective, it might suggest that the indirect effects of unconventional monetary policy in lowering interest rates via bond purchases in the economy might have larger impacts on firms' capital structure than the direct effects of purchasing the bonds that were issued by the private companies.

In regard to differences in unconventional monetary policy effects across various firm sizes, only the debt maturity of medium and large firms are affected. Unconventional policy induced large firms to higher TDTA, LTDTA, and LTDTD by 2.56 bps, 0.811 bps, and 1.32 bps respectively. These results might be because of the unique source of funding for large firms. Larger firms, especially publicly listed firms, have easier access to capital markets, which is an alternative source of funding. The capital market serves as a 'Spare Tire' for large firms to mitigate the risks of shocks due to crisis [Demirguc-Kunt et al. (2015)]. This spare tire hypothesis is also clear in the findings of this research. The crisis shock on LTDTA for medium-sized firms is larger than its effect for large-sized firms. The large firms' reliance on capital markets enlarged the exposure of their capital structure to the market. As unconventional policy impacts capital markets, large firms bore more significant effects of this than any other size of the firm. In the case of unconventional monetary policy, large firms are motivated to take more debt in the short and long term.

Some unconventional monetary policy specifically targets certain sectors as a liquidity injection. The central banks' QE is a program to purchase government bonds. The central banks' quantitative and qualitative easing (QQE) or credit easing (CE) is a program to purchase corporate bonds, ETFs of specific sectors or other particular indices. These purchasing programs target the selected sectors. Although major central banks who employ unconventional policy do not disclose the target sectors, the results indicate unconventional policy affects some sectors but not others. The sectors that have statistically significant UMP impacts on all three capital structure measures are Finance, Producer Manufacturing, and Utilities. The sectors that have statistically significant UMP impacts on both TDTA and LTDTA are Consumer Durables, Consumer Services, Distribution Services, Electronic Technology, Process Industries, and Transportation. The sectors that have statistically significant UMP impacts only on either TDTA or LTDTA are Communications, Consumer Non-Durables and Non-Energy Minerals. Both the country and industry level results discussed above are available on request from the author.

B. Portability of Capital Structure Theories

In order to examine the portability of capital structure theories, the results by country from the corporate finance approach were compared against the findings from past studies. Portability in this context means that the capital structure theories hold across time and through different monetary policy schemes. The estimation from regressions using the corporate finance method select each countries is shown in table 4.

Table 4—Estimation Results by Country

Variables	FATA	ROA	STA	GITA	TA	Log GDP	Constant
Canada							
td/ta	0.0583***	-0.000332***	-0.0396***	-0.00206	-7.76e-07	0.0255***	-0.0379***
	(0.00758)	(7.60e-05)	(0.00463)	(0.0105)	(8.85e-07)	(0.00270)	(0.00307)
ltd/ta	0.0651***	-8.90e-05	0.0132*	0.00170	-7.40e-06***	0.00852***	0.0817***
	(0.0109)	(0.000118)	(0.00673)	(0.0167)	(1.04e-06)	(0.00236)	(0.00547)
ltd/td	0.0220	0.000581	0.00943	0.0186	-3.96e-06	0.0537***	0.204***
	(0.0531)	(0.000708)	(0.0355)	(0.124)	(3.62e-06)	(0.0118)	(0.0698)
France							
td/ta	-0.182***	-0.000885	-0.128***	-0.0623	8.55e-07	0.0454***	-0.0531***
	(0.0382)	(0.00110)	(0.0203)	(0.128)	(1.25e-06)	(0.00476)	(0.00737)
ltd/ta	-0.168**	0.00250	-0.0914*	-0.0302	-1.77e-06	0.0335***	-0.00750
	(0.0724)	(0.00238)	(0.0485)	(0.332)	(2.35e-06)	(0.00877)	(0.0272)
ltd/td	-0.602**	0.0136*	-0.140	0.772	-1.16e-05	0.0892***	0.108
	(0.237)	(0.00794)	(0.173)	(1.254)	(7.70e-06)	(0.0303)	(0.135)
Germany							
td/ta	0.0275***	-0.000933***	-0.0322***	-0.0104	-2.00e-07**	0.0161***	0.0339***
	(0.00951)	(9.15e-05)	(0.00355)	(0.00722)	(8.64e-08)	(0.00204)	(0.00193)
ltd/ta	0.0346***	-0.000471***	0.00267	0.0288**	-3.06e-07***	0.00656***	0.0657***
	(0.0134)	(0.000143)	(0.00535)	(0.0118)	(1.15e-07)	(0.00182)	(0.00365)
ltd/td	0.0274	-3.19e-06	-0.0238	0.168	-5.16e-07	0.0409***	0.278***
	(0.0749)	(0.000991)	(0.0347)	(0.111)	(5.39e-07)	(0.0106)	(0.0585)
Italy							
td/ta	0.0137	-0.00367***	-0.0567***	-0.00673	1.28e-06***	0.0388***	-0.122***
	(0.0123)	(0.000128)	(0.0110)	(0.0204)	(3.34e-07)	(0.00279)	(0.00378)
ltd/ta	0.0647***	-0.000279*	-0.00256	-0.00396	-1.20e-06***	0.0183***	-0.0296***
	(0.0146)	(0.000163)	(0.0133)	(0.0266)	(4.05e-07)	(0.00257)	(0.00558)
ltd/td	0.196***	0.00169*	0.198***	0.0710	-3.90e-06*	0.0270**	0.190***
	(0.0721)	(0.000896)	(0.0684)	(0.154)	(2.08e-06)	(0.0117)	(0.0373)
Japan							
td/ta	0.0277***	-0.00212***	-0.0319***	7.98e-06	-2.86e-08	0.0125***	0.0804***
	(0.00177)	(4.04e-05)	(0.00112)	(6.46e-05)	(5.12e-08)	(0.000817)	(0.000812)
ltd/ta	0.0223***	-0.000734***	0.0117***	-8.38e-05	-2.61e-07***	0.00435***	0.0212***
	(0.00231)	(5.08e-05)	(0.00134)	(8.66e-05)	(6.86e-08)	(0.000497)	(0.00107)
ltd/td	0.0109	-0.000173	0.0292***	-0.000411	-3.56e-07	0.0391***	-0.0538***
	(0.0150)	(0.000300)	(0.00733)	(0.000672)	(4.31e-07)	(0.00220)	(0.00978)
The United Kingdom							
td/ta	0.0143	-0.000701*	-0.165***	-0.00391	-1.84e-06	0.0643***	-0.148***
	(0.0313)	(0.000399)	(0.0273)	(0.0270)	(1.41e-06)	(0.0120)	(0.00656)
ltd/ta	0.0926**	-0.000538	-0.0185	0.0365	-3.27e-06*	0.0268***	-0.0195
	(0.0439)	(0.000626)	(0.0410)	(0.0501)	(1.91e-06)	(0.00786)	(0.0185)
ltd/td	0.216	0.0108***	-0.0196	-0.0961	-2.38e-05***	0.101***	-0.0773
	(0.167)	(0.00246)	(0.166)	(0.240)	(6.80e-06)	(0.0310)	(0.116)
The United States							
td/ta	-0.00326***	-1.09e-05**	0.0346***	0.0588***	7.58e-08	0.0236***	-0.0221***
	(0.000607)	(5.19e-06)	(0.00229)	(0.00492)	(2.68e-07)	(0.00199)	(0.00231)
ltd/ta	-0.00433***	-2.99e-05***	0.0142***	0.0772***	-2.68e-06***	0.0178***	0.0329***
	(0.000407)	(6.92e-06)	(0.00233)	(0.00558)	(2.33e-07)	(0.00125)	(0.00297)
ltd/td	-9.81e-05	6.34e-07	0.0117	0.0692***	-1.41e-06***	0.0684***	0.0613**
	(0.000674)	(3.54e-05)	(0.00726)	(0.0268)	(5.28e-07)	(0.00403)	(0.0259)

Note: *** p<0.01, ** p<0.05, * p<0.1.

Source: FactSet and World Bank Data Repository

Tangibility is measured by fixed assets to total assets (FATA). Although there are exceptions, fixed assets to total assets are positively correlated with TDTA, LTDTA, and LTDTD for the majority of countries. Some countries have negative coefficients on FATA for TDTA. In rare cases, the coefficients on FATA for LTDTA and LTDTA are also negative. Hence, the hypothesis of maturity matching of assets and liabilities are partially portable but not for all countries. Profitability mainly has a negative impact on leverage ratios, especially on TDTA. The profitability is reflected in return on assets (ROA) and gross income to total assets (GITA). The results support the argument from Myers and Majulf and indicate that Pecking-Order Hypothesis (POH) is portable [Myers and Majluf (1984)]. This result also suggests that there is asymmetric information after the global financial crisis and the introduction of unconventional monetary policy. Growth opportunities represented by sales to total assets (STA) are mostly negatively related to capital structure. As Myers points out, profitable firms forego positive net

present value projects [Myers (1977)]. Debt overhang exists in unconventional policy periods as well. Size in terms of total assets (TA) mainly has a negative effect on leverage ratios. This result questions existing theories. The estimated total assets impacts on capital structure do not align with the results from [Demirguc-Kunt et al. (2015)]. These differences might reveal that unconventional monetary policy changed the relationship of size and leverage of firms or affected international competitiveness. The findings indicate that most capital structure theories are applicable during periods of unconventional policy. However, the capital structure and size of the firms does not interact in the same manner ex-ante and ex-post of unconventional policy. This signals a need for future research on the portability of capital structure theories among different monetary policy schemes.

C. The Impact of Conventional vs. Unconventional Policy Shocks on Capital Structure

The estimated effects of conventional and unconventional monetary policy shocks on the capital structure for the horizon of 24 months are shown in figure 1-8. They exhibit the response of variables to the impulse of a one standard deviation increase in the orthogonal structural shock u_t^z in equation 9. The transmission of monetary policy shocks to the capital structure measures vary amongst countries and across conventional and unconventional periods. The types of unconventional policy employed shows clear cross-country variation. The unconventional policy shocks seem to have more gradual, consistent and persistent effects on corporate leverage than the conventional policy shocks. The discussion here is centred on the response of capital structure measures to the impulse response of monetary policy.

The figures indicate the responses of TDTA, LTDTA and LTDTD to a monetary shock for 24 months after the shock. The solid navy line represents the median response. The 95% confidence interval bands (i.e. $\pm$ two standard deviations from the median response) of the response is shown in dotted blue lines. The first and second row show the impulse responses during conventional and unconventional periods, respectively. The first column is the response of TDTA to the impulse of the monetary shock. The second column is the response of LTDTA. The third column is the response of LTDTD. TDTA is total debt to total assets. LTDTA is long-term debt to total assets. LTDTD is long-term debt to total debt. TDTA, LTDTA and LTDTD, are the national average for each month. These national averages are computed after the restrictions of TDTA, LTDTA and LTDTD<10.00 are imposed. The impulse of monetary shock is the one standard deviation increase in the structural shock of the monetary policy indicator equation in SVAR. The vertical axis indicates the percentage change of each responding variable, while the horizontal axis is the months after the shock. Results are reviewed in the order of Canada, France, Germany, Italy, Japan, Sweden, Switzerland, the United Kingdom, and then the United States. [13]

[13]For Canada and the US, the conventional periods are from 2001M1 to 2008M9 and the unconventional periods are from 2008M10 to 2018M12. For France, Germany, and Italy, the conventional periods are from 2001M1 to 2015M1 and the unconventional periods are from 2015M2 to 2018M12. For Japan, the conventional periods are from 2001M1 to 2010M6 and the unconventional periods are from 2010M7 to 2018M12. For Sweden, the conventional periods are from 2001M1 to 2014M12 and the unconventional periods are from 2015M1 to 2018M12. For Switzerland and the UK, the conventional periods are from 2001M1 to 2008M12 and the unconventional periods are from 2009M1 to 2018M12.

Figure 1. Impulse Responses in Canada

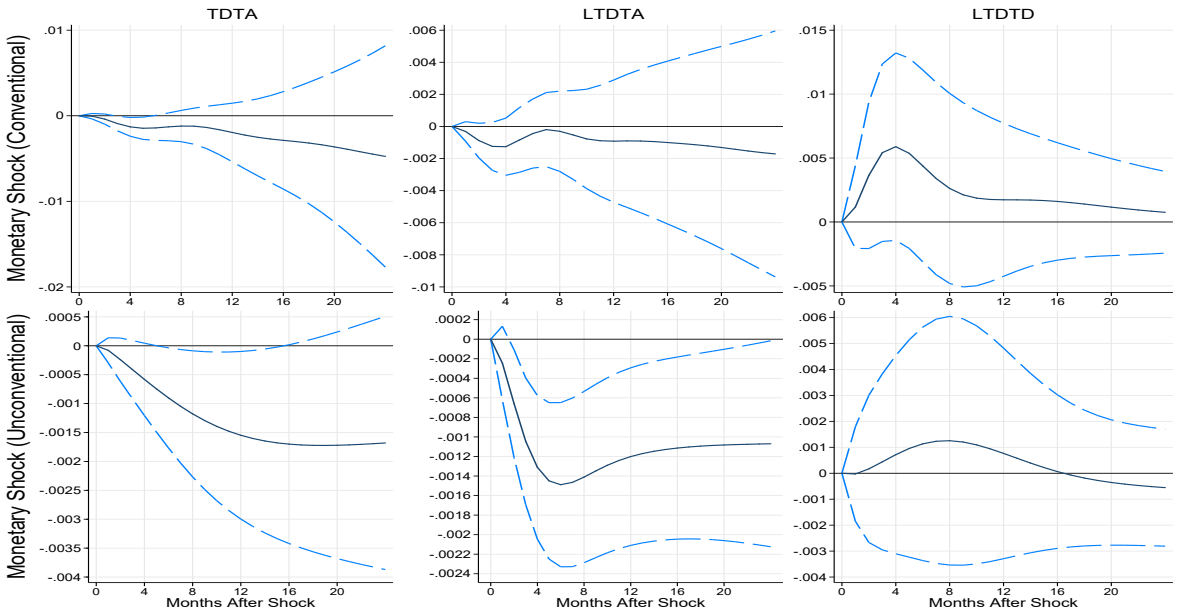

The Bank of Canada implemented unconventional monetary policies in April 2009. The impulse responses in figure 1 indicate different movements ex-ante and ex-post of the introduction of unconventional policy. The first row in figure 1 shows the impulse response for the conventional period and the second row for the unconventional period. During the conventional period, monetary policy tightening decreases TDTA and LTDTA gradually over the span of 24 months, whereas during the unconventional periods, TDTA and LTDTA react to the impulse much quicker for the first 12 months. The response of LTDTD after 16 months from the shock demonstrates the opposite response for conventional and unconventional periods. Following the global financial crisis, firms deleveraged their long-term debt and lowered debt maturity [Demirguc-Kunt et al. (2015)]. The debt maturity measure, LTDTA, rises within the first quarter after the monetary shock. The increase in LTDTD is larger for conventional periods than for unconventional periods. Both TDTA and LTDTA decline for the initial four months after the shock then start rising. The responses of all capital structure measures indicate that the interactions of monetary policy and capital structure differ across different policy schemes, conventional and unconventional. Compared with other countries, the responses of leverage ratios are relatively moderate in Canada.

Figure 2. Impulse Responses in France

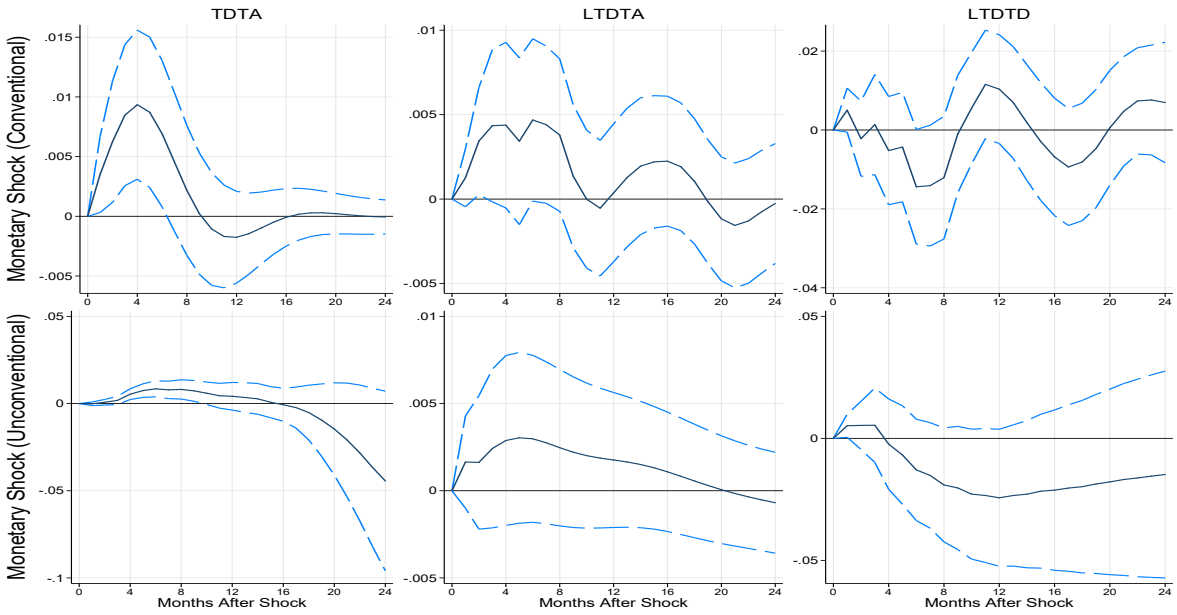

The European Central Bank conducted forward guidance in 2013 followed by the introduction of negative interest rate in 2014 and large-scale asset purchase programme in 2015 [Hartmann and Smets (2018)]; [Constncio (2018)]. For France, the capital structure responds to conventional and unconventional monetary policy differently. While there are fluctuations in firms financing decisions after a conventional policy shock, firms tend to respond more smoothly for the unconventional policy shock. For TDTA, the ratio slowly increases for the first 5 months, then decreases at an increasing rate in the unconventional case. TDTA seems to have no response 16 months after a conventional monetary policy shock. For LTDTA, companies adjust their long-term debt multiple times within 2 years after a contractionary conventional policy. The impulse response of unconventional monetary tightening on LTDTA is absorbed within 20 months. The enterprises initially expand their long-term debt by 0.25 bps for a quarter, then gradually slow their expansion of LTDTA. The debt maturity, does not chase the LTDTA response. LTDTD rises for the first 4 months by 0.1 bps then reduces by 0.25 bps in unconventional periods. After a year from an unconventional monetary policy shock, the reduction of the firm's debt maturity slows down. The relatively more stable responses in the unconventional period might have resulted from forward guidance (FG). Once the future path of the countries' monetary policy is described in FG, firms might adjust their behaviour accordingly. The response of TDTA and LTDTA for conventional periods conflicts with the economic theory. It is believed that contractionary monetary policy discourages enterprises to acquire additional debt.

Figure 3. Impulse Responses in Italy

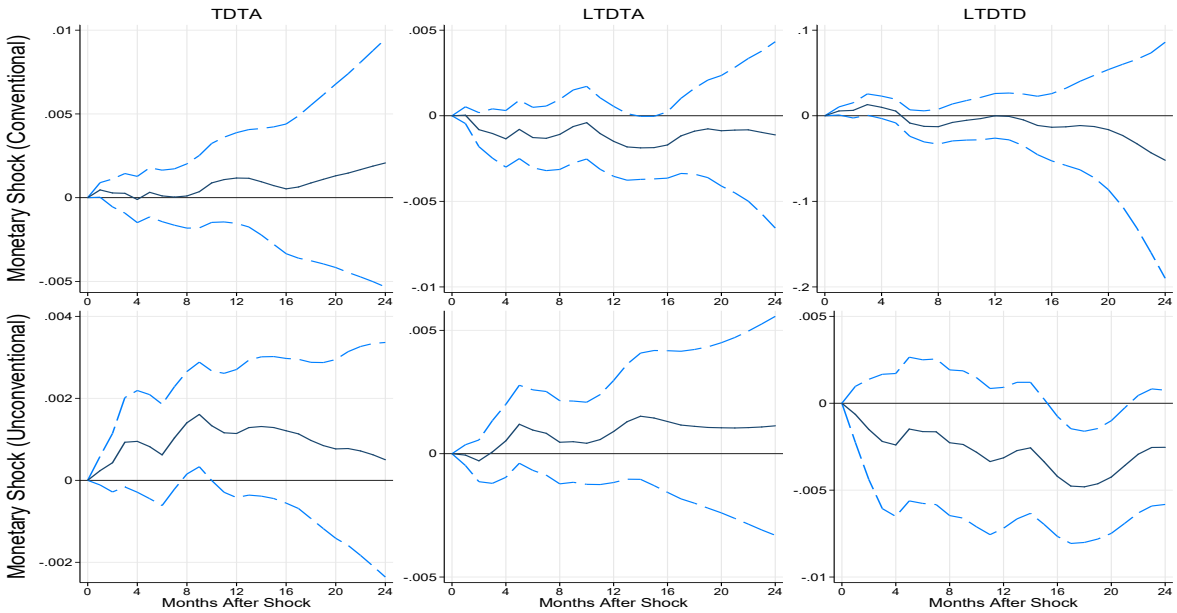

For Italy, the patterns of leverage responses are analogous across the monetary policy schemes. The magnitude of responses tend to be larger during the conventional period than in the unconventional period. There are divergences in the responses of TDTA and LTDTD after 16 months. TDTA rises relatively faster after 16 months from the surprise of conventional policy tightening. However, after 16 months from an unconventional policy shock, the change in TDTA decreases indicating that firms slow down the pace of taking more debt. Similarly, deleveraging LTDTD accelerates after 20 months from a conventional policy shock, whilst deleveraging LTDTD decelerates 16 months after the contractionary unconventional policy shock. LTDTA, on the other hand, decreases from conventional policy shocks and increases from unconventional policy shocks.

Germany, shares the same monetary policy with France and Italy, hence its figures have been omitted. One point of divergence is the relationship of leverage response to the impact of monetary policy. During the two years after the shock, all three capital structure measures react in the opposite direction. During the conventional periods, TDTA, LTDTA, and LTDTD increase by 0.125 bps, 0.18 bps, and 0.2 bps at their peak, respectively. During the unconventional period, changes in TDTA and LTDTA fluctuate between 0.7 bps, plus or minus, for both leverage ratios. These fluctuations are also reflected in LTDTD. LTD increases by 0.35 bps roughly for three months then decreases by 0.2 bps for 16 months.

Figure 4. Impulse Responses in Japan

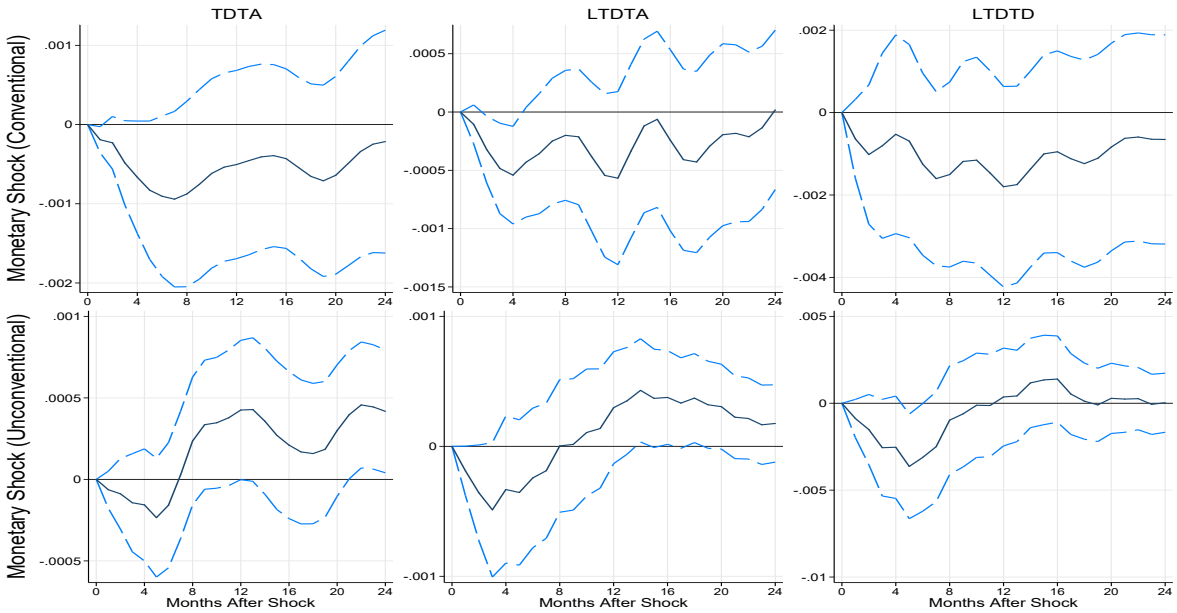

Japan was the front-runner in implementing unconventional monetary policy. The Bank of Japan (BoJ) implemented an unconventional-like policy in 2006 as the country was already struggling with low-interest rates from the "Lost Decades". The response to the global financial crisis was relatively minimal as the BoJ conducted forward guidance and a few asset purchases from 2010 to 2012. From 2013, the BoJ started to employ unconventional policy extensively with a massive scale of asset purchases and introduced yield curve controls (YYC) with negative interest rates (NIR) after central bank reserves. The capital structure responses to CMP shock follow classical economic theory. By tightening the policy rates, firms reduce their debt. The monetary shocks almost fade away in 2 years. Changes in TDTA and LTDTD seem to converge to 0 after 2 years from the shock. On the other hand, the capital structure measures increased 6-12 months after the unconventional policy shock. TDTA declined for the first 7 months by 0.02 bps at the peak then started rising by 0.05 bps. The decrease in LTDTA diminished in 8 months and then continued to rise afterward. The response of LTDTD fluctuates over 20 months.

Figure 5. Impulse Responses in Sweden

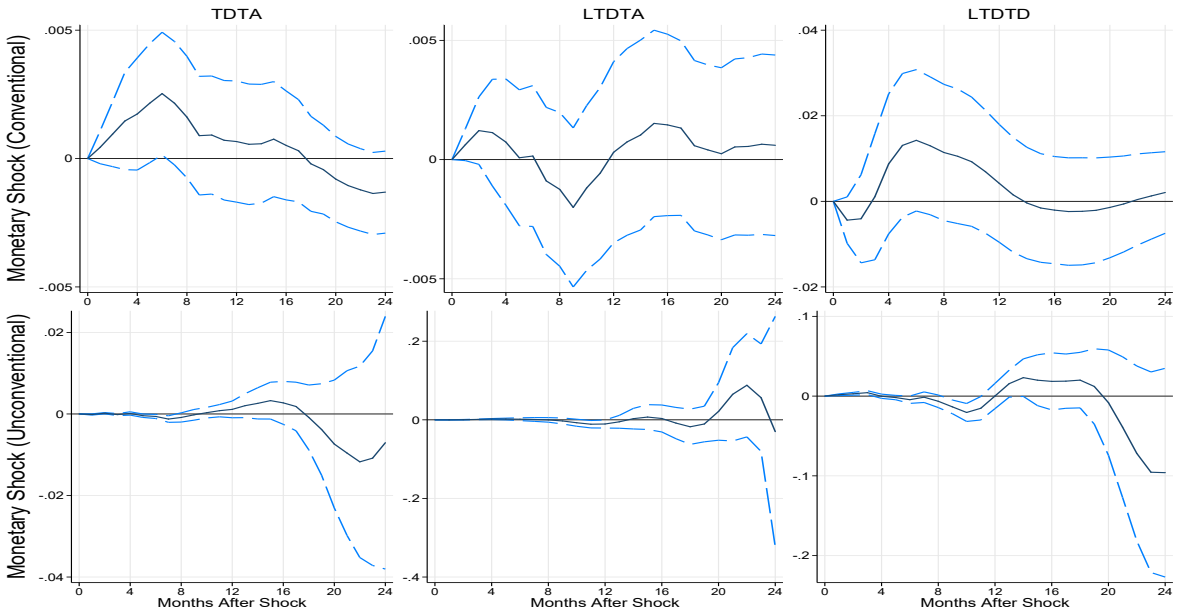

Riksbank, the central bank of Sweden, officially switched their monetary policy from conventional to unconventional in February 2015 [Rezende (2017)]. Riksbank's unconventional monetary policy consists of three programs: forward guidance, quantitative easing and a negative interest rate policy. In the conventional period, TDTA increases in 6 months by 0.25 bps but decreases in 2 years by 0.125 bps. The response from a conventional policy shock fluctuates for LTDTA around ± 0.2 bps and for LTDTD between -0.5 bps and +18 bps. There were no notable responses in capital structure until 8 months after the unconventional policy shock. The dynamism of policy transmission changes with unconventional monetary policy. TDTA reveals a 10 bps decline in almost 2 years after the impulse of unconventional monetary policy. Firms increase their LTDTA by 8 bps in 2 years after UMP shock. LTDTD fluctuates ± 2bps for 14 months after the UMP shock then declines by 10 bps in 2 years. The responses of capital structure measures are significantly larger for unconventional than for conventional policy shocks.

FIGURE 6. IMPULSE RESPONSES IN SWITZERLAND

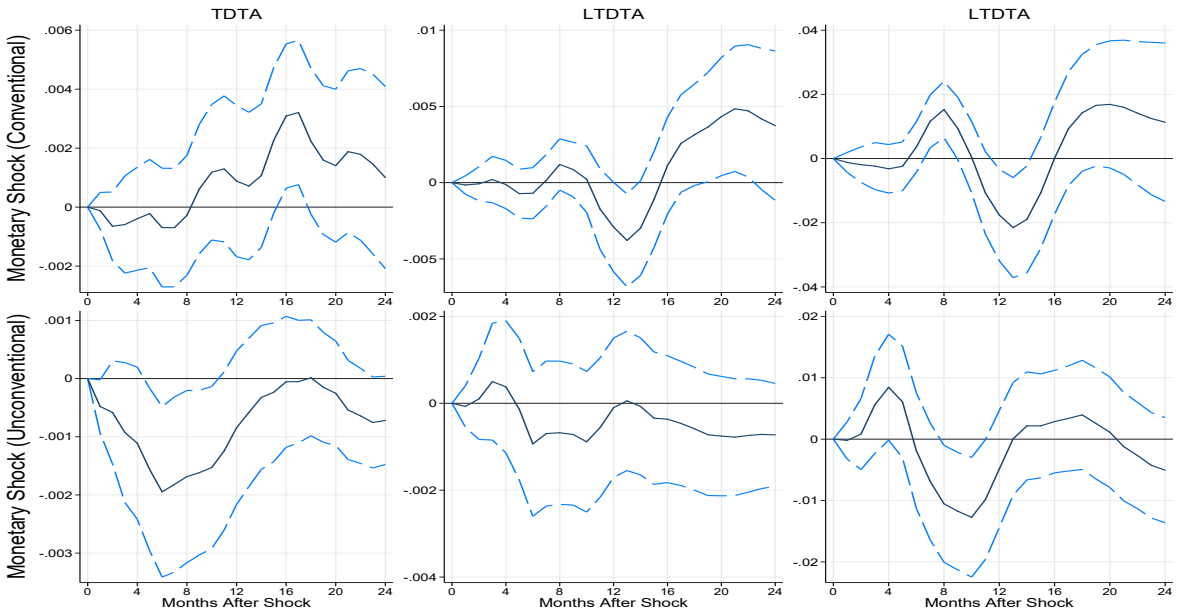

The Swiss National Bank (SNB) announced the introduction of unconventional monetary policy, in September 2011. However, since March 2009, the SNB intervened in foreign exchange markets and commenced a massive-scale foreign currency purchase [Maechler (2016)]. The SNB kept the minimum exchange rate set against the Euro until January 2015 [Jordan (2016)]. At the beginning of 2015, the SNB started charging a negative interest rate of 0.75% to on-site deposits at the central bank. This paper uses the date that the SNB initiated a large-scale foreign currency purchase as the start of the unconventional period. The response of firms' leverage to conventional and unconventional policy shocks follows similar patterns but their magnitude varies across policy schemes. The conventional policy shocks are generally larger than the unconventional policy shocks. TDTA levels decline but go back to pre-shock levels for conventional monetary tightening. However, TDTA continues to decrease even two years after an unconventional policy shock. This is evidence of the long-lasting effect of unconventional monetary policy on TDTA. The same is true for LTDTA, which does not return to the pre-shock level even 2 years after the unconventional policy shock. The effect of monetary policy shocks on a firms capital structure is smaller in the short run but more persistent for unconventional monetary policy.

Figure 7. Impulse Responses in the UK

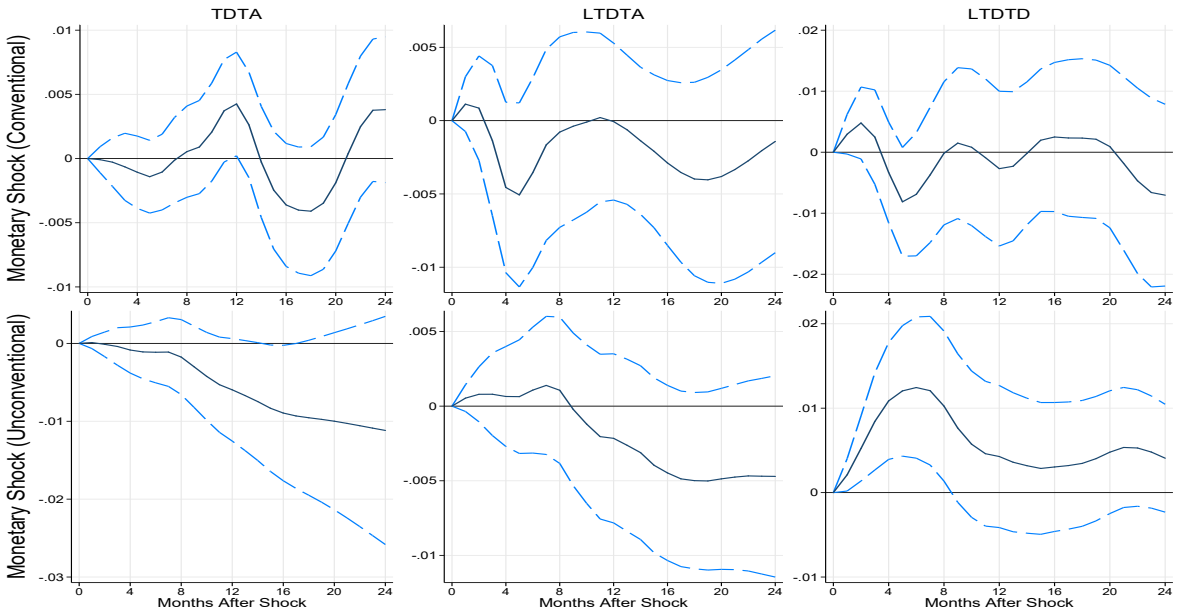

The literature suggests there were three phases of unconventional monetary policy in the UK [Dell'Ariccia et al. (2018)]. The first stage was the establishment of large-scale quantitative easing programs between 2009 and 2012 to recover from the global financial crisis. The second stage was forward guidance to commit no further policy tightening for 2013-2014. Lastly, there was another round of QE due to concerns about "Brexit." The responses of capital measures were larger and less volatile for unconventional policy shocks than for conventional policy shocks. While there were some recoveries of leverage ratios from the unexpected shock in conventional periods, leverage responses in unconventional periods do not indicate any recovery except in LTDTD. Analogous to Switzerland, unconventional monetary policy effects are long-lasting for both TDTA and LTDTA. The forward guidance signals the future policy rate, hence, there were fewer fluctuations in both TDTA and LTDTA for the unconventional period. The QE might be successful in providing liquidity, but the impulse-response functions during UMP indicate that QE does not act as a buffer to mitigate the impact on firms financing decisions.

FIGURE 8. IMPULSE RESPONSES IN THE US

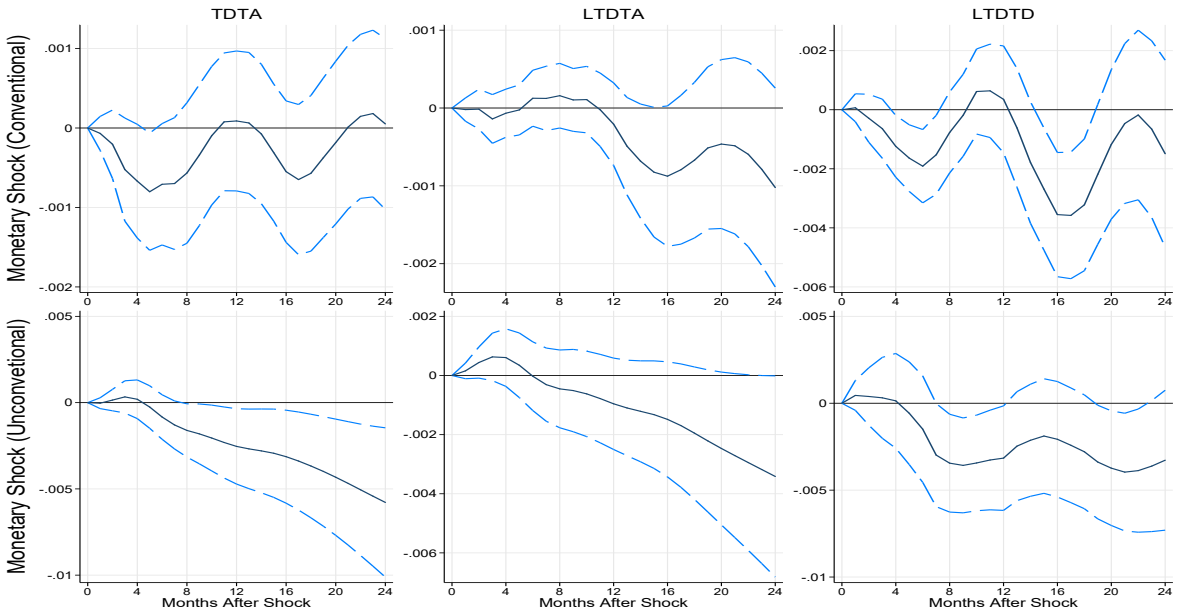

The Federal Reserve (Fed) reacted to the global financial crisis quicker than any other central bank. The Fed started to implement unconventional monetary policy including quantitative easing and forward guidance in November 2008. The Fed terminated its unconventional policy in late 2014 and since October 2017 their balance sheet has been shrinking as existing securities matured [Kuttner (2018)]. The responses of capital structure for unconventional periods are somewhat similar to the case of the UK. The impulse response may indicate that these two countries experienced analogous unconventional monetary policy effects on firms' capital structure. This result seems reasonable as the US and the UK employed identical unconventional policies and share many common country/institutional factors. In unconventional periods there was a consistent deleveraging of TDTA and LTDTA. LTDTD demonstrates a moderate recovery for certain periods, but the ratio is near 0.5 bps lower than the pre-shock level by the second year following the shock. For unconventional monetary policy shocks, companies in the US were not able to recover from the seemingly persistent unconventional policy effects on their capital structure.

The response of capital structure measures to the conventional and unconventional policy shocks are different in every country. The results indicate that there is a structural change in an economy after unconventional policy is introduced. The cross-country variations in the responses of firm leverage to an unconventional policy shock show that the effects and transmission mechanisms vary depending on the types of unconventional monetary policy (ex. FG,

QE and NIR etc.) and on the country's economic conditions. The gradual, consistent, and persistent responses of capital structure measures to unconventional policy shocks compared to conventional policy shocks might indicate that the intended reactions can be obtained by the successful implementation of unconventional monetary policy. The different impulse responses amongst France, Germany, and Italy despite sharing identical monetary policy shows that country characteristics play an important role in disseminating monetary policy shocks on firm leverage ratios. The analogous patterns of impulse responses during unconventional periods between the UK and the US, albeit with different monetary policies and country characteristics, suggest the possibility that the reaction of firm financing decisions to monetary policies can be managed with unconventional monetary policy.

VI. Conclusion

A rise in the use of unconventional monetary policy changed the dynamic relationship between monetary policy and firm financing decisions. To study this, empirical settings from corporate finance and monetary economics were applied to the firm- and country-level data. The corporate finance approach revealed variation in the effect of unconventional monetary policy on capital structure. This variation can be partially explained by country-, sector-, and firm-specific factors. The unconventional policy effects were significant for large private companies in upper-middle-income countries that did not experience the global financial crisis. Furthermore, classical capital structure theories were applicable and portable across policy schemes, conventional and unconventional. By using SVAR it was determined that the transmission of monetary policy shocks on firm capital structure is different amongst conventional and unconventional periods for each country. There are country specific factors that affect the transmission mechanisms of unconventional monetary policy in the economy. The common attributes in the response of corporate leverage ratios to an unconventional monetary policy shock might reveal that the effects of monetary policy on firm capital structure is better managed with unconventional monetary than conventional monetary policy. Answers are now available for the three questions posed in the introduction: firstly it is clear that the unconventional monetary policy induced higher corporate leverages after the global financial crisis; secondly most of the capital structure theories are still applicable; and thirdly the reactions of firms' financing decisions to unconventional policy shocks are more moderate, constant and long-lasting than to conventional policy shocks.

REFERENCES

[Akaike 1998] Akaike, H: *Information Theory and an Extension of the Maximum Likelihood Principle*. NY: Springer, 1998

[Alsterlind et al. 2015] Alsterlind, J. ; Sandstrom, M. ; Vestin, D.: How can government bond purchases make monetary policy more expansive? In: *Economic Commentraries, Sveriges Riksbank* 12 (2015), S. 1–8

[Amiti and Weinstein 2013] Amiti, M. ; Weinstein, D.: How much do idiosyncratic bank shocks affect investment? Evidence from matched bank-firm loan data. In: *NBER Working Papers* No. 18890 (2013). – URL `https://doi.org/10.3386/w18890`

[Bernanke 2014] Bernanke, B.: The Fed yesterday, today and tomorrow. (2014)

[Bernanke and Gertler 1995] Bernanke, B. ; Gertler, M.: Inside the Black Box: The Credit Channel of Monetary Policy Transmission. In: *NBER Working Papers* No. 5146 (1995). – URL https://doi.org/10.3386/w5146

[Booth et al. 2001] Booth, L. ; Aivazian, V. ; A., Demirguc-Kunt ; Maksimovic, V.: Capital Structures in Developing Countries. In: *The Journal of Finance* 56 (2001), Nr. 1, S. 87–130. – URL https://doi.org/10.1111/0022-1082.00320

[Cahn et al. 2017] Cahn, C. ; Mullins, W. ; Duquerroy, A.: Unconventional Monetary Policy and Bank Lending Relationships. In: *Bank of France Working Papers* No. 659 (2017). – URL https://doi.org/10.2139/ssrn.2970199

[Campbell et al. 2012] Campbell, J. ; Evans, C. ; Fisher, J. ; Justiniano, A.: Macroeconomic Effects of Federal Reserve Forward Guidance. (2012). – URL https://doi.org/10.1353/eca.2012.0004

[Constncio 2018] Constncio, Vtor: *Past and future of the ECB monetary policy.* 5 2018

[Curcuru et al. 2018] Curcuru, S. ; Kamin, S. ; Li, C. ; Rodriguez, M.: International Spillovers of Monetary Policy: Conventional Policy vs. Quantitative Easing. In: *FRB International Discussion Papers* No. 1234 (2018). – URL https://doi.org/10.17016/ifdp.2018.1234

[Dell'Ariccia et al. 2018] Dell'Ariccia, G. ; Rabanal, P. ; Sandri, D.: Unconventional monetary policies in the Euro Area, Japan, and the United Kingdom. In: *Journal of Economic Perspectives* 32 (2018), Nr. 4, S. 147–172. – URL https://doi.org/10.1257/jep.32.4.147

[Demirguc-Kunt et al. 2015] Demirguc-Kunt, A. ; Martinez-Peria, M. ; Tressel, T.: The impact of the global financial crisis on firms' capital structure. In: *World Bank Policy Research Working Papers* No. 7522 (2015). – URL https://doi.org/10.1596/1813-9450-7522

[Fawley and Juvenal 2012] Fawley, B. ; Juvenal, L.: Quantitative easing: lessons we've learned. In: *The Regional Economist, Federal Reserve Bank of St. Louis* 7 (2012), S. 8–9

[Fiedler et al. 2016] Fiedler, S. ; Jannsen, N. ; Wolters, M. ; Hanisch, I.: Transmission channels of unconventional monetary policy in the euro area: where do we stand? In: *EuroParl* (2016). – URL

[Gertler and Karadi 2011] Gertler, M. ; Karadi, P.: A model of unconventional monetary policy. In: *Journal of Monetary Economics* 58 (2011), Nr. 1, S. 17–34. – URL https://doi.org/10.1016/j.jmoneco.2010.10.004

[Gertler and Karadi 2015] Gertler, M. ; Karadi, P.: Monetary policy surprises, credit costs, and economic activity. In: *American Economic Journal: Macroeconomics* 7 (2015), Nr. 1, S. 44–76. – URL https://doi.org/10.1257/mac.20130329

[Haldane et al. 2016] Haldane, A. ; Robert-Sklar, M. ; Wieladek, T. ; Young, C.: QE: the story so far. In: *Bank of England Working Papers* No. 624 (2016). – URL https://www.bankofengland.co.uk/working-paper/2016/qe-the-story-so-far

[Harris and Raviv 1991] Harris, M. ; Raviv, A.: The Theory of Capital Structure. In: *The Journal of Finance* 46 (1991), Nr. 1, S. 297–355. – URL https://doi.org/10.1111/j.1540-6261.1991.tb03753.x

[Hartmann and Smets 2018] Hartmann, Philipp ; Smets, Frank: The first twenty years of the European Central Bank: monetary policy. (2018), Nr. 2219. – URL https://ideas.repec.org/p/ecb/ecbwps/20182219.html

[Jordan 2016] Jordan, T.: *Monetary policy using negative interest rates: a status report.* 2016

[Kashyap and Stein 1994] Kashyap, A. ; Stein, J.: The Impact of Monetary Policy on Bank Balance Sheets. In: *NBER Working Papers* No. 4821 (1994). – URL https://doi.org/10.3386/w4821

[Kuttner 2018] Kuttner, K.: Outside the box: unconventional monetary policy in the Great Recession and beyond. In: *Journal of Economic Perspectives* 32 (2018), Nr. 4, S. 121–146. – URL https://doi.org/10.1257/jep.32.4.121

[Laeven and Valencia 2018] Laeven, L. ; Valencia, F.: Systemic banking crises revisited. In: *IMF Working Papers* No. 18/206 (2018). – URL https://doi.org/10.5089/9781484376379.001

[Lhuissier and Szczerbowicz 2018] Lhuissier, S. ; Szczerbowicz, U.: Monetary policy and corporate debt structure. In: *Bank of France Working Papers* No. 697 (2018). – URL https://doi.org/10.2139/ssrn.3267904

[Maechler 2016] Maechler, A.: Financial markets and monetary policy implementation - an evolving relationship. Transcript (2016)

[Mokhova and Zinecker 2014] Mokhova, K. ; Zinecker, M.: Macroeconomic factors and corporate capital structure. In: *Procedia - Social and Behavioral Sciences* 110 (2014), S. 530–540. – URL https://doi.org/10.1016/j.sbspro.2013.12.897

[Myers 1977] Myers, S.: Determinants of corporate borrowing. In: *Journal of Financial Economics* 5 (1977), Nr. 2, S. 147–175. – URL https://doi.org/10.1016/0304-405x(77)90015-0

[Myers 1984] Myers, S.: The Capital Structure Puzzle. In: *The Journal of Finance* 39 (1984), Nr. 3, S. 574–592. – URL https://doi.org/10.1111/j.1540-6261.1984.tb03646.x

[Myers and Majluf 1984] Myers, S. ; Majluf, N.: Corporate financing and investment decisions when firms have information that investors do not have. In: *Journal of Financial Economics* 13 (1984), Nr. 2, S. 187–221. – URL https://doi.org/10.1016/0304-405x(84)90023-0

[Rajan and Zingales 1995] Rajan, R. ; Zingales, L.: What do we know about capital structure? Some evidence from international data. In: *The Journal of Finance* 50 (1995), Nr. 5, S. 1421. – URL https://doi.org/10.2307/2329322

[Rezende 2017] Rezende, R.: The interest rate effects of government bond purchases away from the lower bound. In: *The Journal of International Money and Finance* 74 (2017), S. 165–186. – URL https://doi.org/10.1016/j.jimonfin.2017.03.005

[Schwarz 1978] Schwarz, G.: Estimating the Dimension of a Model. In: *The Annals of Statistics* 6 (1978), Nr. 2, S. 461–464. – URL https://www.jstor.org/stable/2958889

[Series 2017] Series, The International Monetary Fund P.: Negative Interest Rate Policies. In: *IMF Policy Papers* No. 3 (2017)

[Sims 1980] Sims, C.: Macroeconomics and reality. In: *Econometrica* 48 (1980), Nr. 1, S. 1–48. – URL https://doi.org/10.2307/1912017

[Stock and Watson 2001] Stock, J. ; Watson, M.: Vector Autoregressions. In: *Journal of Economic Perspectives* 15 (2001), Nr. 4, S. 101–115. – URL https://doi.org/10.1257/jep.15.4.101

APPENDIX

TABLE A1—DATA SUMMARY

	Without Restrictions			With Restrictions		
	Freq.	Percent	Cum.	Freq.	Percent	Cum.
With/ Without GFC shock						
No GFC shock	1,037,200	74.58	74.58	372,378	77.11	77.11
GFC shock	353,600	25.42	100	110,518	22.89	100
Income Classes						
Low	10,530	0.75	0.75	779	0.16	0.16
Lower Middle	201,771	14.33	15.08	47,742	9.89	10.05
Upper Middle	360,045	25.57	40.64	129,088	26.73	36.78
High	835,839	59.36	100	305,287	63.22	100
Company Types						
Holding Company	8,100	0.58	0.58	1,326	0.27	0.27
Joint Venture	567	0.04	0.62	48	0.01	0.28
Non Profit Organization	2,349	0.17	0.78	212	0.04	0.33
Private Company	76,707	5.45	6.23	2,807	0.58	0.91
Public Company	1,253,718	89.03	95.26	463,874	96.06	96.97
Subsidiary	66,744	4.74	100	14,629	3.03	100
Company Sizes						
Small	161,676	11.48	11.48	28,928	5.99	5.99
Medium	159,894	11.35	22.84	38,977	8.07	14.06
Large	1,086,615	77.16	100	414,991	85.94	100

Source: FactSet and World Bank Data Repository

Table A2—Lag Selections

	TDTA		LTDTA		LTDTD	
	Con	Unconv	Con	Unconv	Con	Unconv
Canada	4	2	4	2	2	2
France	8	8	8	8	7	8
Germany	3	8	3	8	7	8
Italy	8	8	8	8	8	8
Japan	8	8	8	6	8	6
Sweden	8	8	8	8	5	8
Switzerland	8	8	8	8	8	8
The United Kingdom	8	8	5	8	8	8
The United States	8	5	8	5	8	7

Note: Lags are selected according to Akaike's and Bayesian information criteria. TDTA is total debt to total assets. LTDTA is long-term debt to total assets. LTDTD is long-term debt to total debt. Con indicates conventional periods. Unconv refers to unconventional periods. Each number represents the number of monthly lags.

Modeling Social Learning Using Dyna-Q and Ant Colony Optimization

By Renato Zimmermann

This paper introduces a novel way of modeling social learning in macroeconomics using techniques from model-based reinforcement learning and ant colony optimization. The work extends previous works in bounded rationality and social learning by providing tools to complement previously-distinct models in adaptive learning. We test these new techniques using simulations of job search and consumption. Results demonstrate that models fit using the proposed techniques can learn core economic behaviours given no information about the environment, but do not fully fit reward functions in line with rational expectations theory.

***Keywords**: Optimization, Machine Learning, Dyna, Learning Algorithms, Adaptive Learning, Job Search, Information Economics*

I. Introduction

The way humans perceive their environment dictates how they make decisions; this perception is ever-changing, based on previous perceptions and influenced by personal experience. Transitioning from one belief to another is often volatile and subject to the whims of those who influence public information the most. Work in cognitive psychology and behavioural economics have shown that consistent biases in decision-making are possible and that herd behaviour on the part of individuals can lead to irrational decision making [Tversky and Kahneman (1982)] [Lin et al. (2013)]. This paper aims to provide a computational framework on which to model information transmission between agents in social settings using minimal assumptions of agent rationality or knowledge.

The rational expectations (RE) framework dominated macroeconomic models during late twentieth century, and to a large extent remains the predominant technique used in modern models [Muth (1961)]. The theory came as an improvement over earlier work using adaptive expectations, which modeled agents' forecasting as a weighted moving average over past values of the predicted variable [Fisher (1911)]. Unlike adaptive expectations, rational expectations theory treats agents' aggregate predictions to be as good as the actual expected value of what they are trying to predict plus some idiosyncratic error. In light of the empirical results that brought RE models to prominence, the field of bounded rationality has aimed to reconcile these models with the possibility of irrational decision making amongst agents. Adaptive learning (AL) is a popular approach among these models, treating agents as non-perfect — but possibly unbiased — forecasters [Evans and McGough (2020)].

This paper aims to extend the branch of AL known as *social learning*, first developed by Arifovic Arifovic (1994). Social learning aims to model how one agent learns in the context of other agents' decisions as well as their own. We contribute to the literature by using the Dyna-Q model-based architecture [Sutton (1991)]. In this respect we represent societal beliefs with a distributed fitting approach inspired by Ant Colony Optimization [Dorigo et al. (2006)]. The techniques proposed here also enable models that connect both social and individual learning, unlike previous techniques that relied on one approach.

We also explore how individual learning can take place without full knowledge of the environment's dynamics and rewards through approximated dynamic programming. Although exact dynamic programming methods yield quick and unbiased solutions, the procedure's transition from

one state of belief to another is not informative in and of itself. The methods outlined in this paper maintain key properties of commonly used dynamic programming techniques while introducing decentralized information dynamics based on socially-learned models and possible channels for misinformation. This work will enable a more realistic study of how autonomous agents influence the dissemination of information based on experience in an environment with unknown dynamics. Overall this paper aims to provide an intuitive framework to model both social and individual learning in tandem, while keeping assumptions about agents' knowledge of underlying stochastic processes to a minimum. This paper is accompanied by an open-source library capable of replicating the results presented here and written with the objective of being easily extended. This framework can serve to facilitate further exploration of social learning and agent-based modeling in economic research.

II. Related Work & Motivation

A. Economic Modelling

One of the first formulations of agent-level forecasting in macroeconomic models is attributed to the work of Irving Fisher in the form of adaptive expectations [Fisher (1911)]. A simple formulation treats agents' forecasting as an integrated moving average IMA(1,1) prediction. Value expectations are simply the previous forecast corrected by its error weighted by some constant coefficient. Using the Cobweb model as a simple example, agents would forecast prices to be:

$$p_t^e = p_{t-1}^e + \lambda(p_t - p_{t-1}^e) \tag{1}$$

where we denote p_t^e, as the expected price forecast at time t. This is solved given p_t, the actual price at time t, and λ a constant coefficient. This simple formulation served as the backbone to a large body of work in the early twentieth century, yet representing forecasting in such a simple manner allowed for counterintuitive behaviour on the part of agents. For example, as the "error correction" term in equation 1 is constant, agents could consistently underestimate prices that trend upwards. This goes against the intuition that agents are rational in nature, and capable of comprehending simple trends.

The rational expectations hypothesis was first posed by John Muth in an effort to reconcile the notion that agents are, in aggregate, rational in nature and capable of forward-looking forecasting [Muth (1961)]. Instead of forecasting values solely using past observations, as was done in adaptive expectations, agents in RE models take into account all available information about a random variable. Their forecasts are the actual expected value of said variable and only wrong by random idiosyncratic noise[1]. Continuing our example using the Cobweb model, price expectations for a representative agent in the rational expectations model would become:

$$p_t^e = \mathbb{E}_{t-1}[p_t] \tag{2}$$

where $\mathbb{E}_{t-1}[p_t]$ denotes the expected value of price at period t given information available at period $t-1$. This formulation allows for a forward, rather than backward, view of agent predictions.

Work on rational expectations was further developed and popularized by Lucas and now serves as a basis for modelling decision making in macroeconomic models [Lucas (1981)]. Regardless, this

[1]It is important to emphasize that the model does not assume that *individual* agents have this knowledge, but rather that their aggregate predictions would be the same as that of a representative agent fully aware of the distribution and parameters of a random process

framework might not fully conform to the realities of how decisions are made in various situations. The assumption that aggregate forward-looking decisions are equal to the true expected future value is arguably too confident of the aggregate capability of agents, as it would imply knowledge of exact distributions and parameters. Granted this argument is sound when dealing with variables of broad interest, it is less convincing for situations with limited price feedback, unsophisticated agents or information frictions.

In order to tackle this issue, the study of bounded rationality has shown that it is possible to limit the capacity of agents' predictions while preserving key theoretical results from rational expectations theory. The adaptive learning (AL) approach has grown to prominence modeling boundedly-rational agents. The approach is based on the *cognitive consistency principle*, which states that agents are just as good at forecasting as economists [Evans and McGough (2020)]. Instead of setting agent predictions to a value's conditional expectation as in RE, Adaptive Learning allows for misspecification of models and other scenarios on the path towards equilibrium. Consider the common implementation of the AL framework known as least-squares learning. As the name suggests, predictions are made through a least-squares estimate on observable variables. The price prediction from before becomes:

$$p_t^e = \alpha'_{t-1} + \beta'_{t-1} w_{t-1} \tag{3}$$

where w_{t-1} is a vector of observable variables in period $t-1$ and α' and β' are least-squares estimates of the perceived model:

$$p_t = \alpha + \beta w_{t-1} + \varepsilon$$

Predictions done this way bridge the backward-looking view from adaptive expectations while letting agents act rationally, yet bounded in knowledge. Another implementation of AL formulated by Jasmina Arifovic uses the concept of social learning to guide how agents make decisions [Arifovic (1994)]. The original social learning technique consists of using a genetic algorithm (GA) to model the evolution of agent's decision making through adaptation and market selection pressures. The algorithm is structured to resemble the process of genetic mutation and evolution that happens in the natural world. One key assumption backing the use of GA is that analogous processes happen in economic environments through a market selection process. The GA approach is unlike those previously outlined in this section, as it specifies agent decisions directly, instead of deriving them as a function taking in some expected value as an input. In the Cobweb model example used thus far, while previous approaches would first generate a price prediction used as an input to a decision function, the GA approach produces this decision directly by encoding the agent's genes. Consider a production function as an example. While the previous approaches would use expected price p_t^e as the input to a production function the GA approach would model the production decision directly as in equation 4:

$$q_{i,t} = \frac{1}{\bar{K}} \sum_{k=1} a_{i,t}^k 2^{k-1} \tag{4}$$

where $q_{i,t}$ is agent i's production at time t and $\bar{K}$ is a normalizing parameter. Genes in the original social learning paper are strictly binary; in equation 4, firm i's k^{th} gene at time t is represented by $a_{i,t}^k$. This approach treats decision-making as an evolving problem, where genes that inform decisions are selected based on which combination yields the best production strategy over time.

Using a genetic algorithm as presented above still lacks a detailed treatment of how agents could use other information available to them, or how actual learning can take place beyond the agent's "genes".

Unlike previous work in adaptive learning, Arifovic's social learning model approaches the problem through an agent based modeling (ABM) standpoint. Developments applying ABM in the context of macroeconomics have been sparse [Farmer and Foley (2009)]. Nonetheless, ABM applications in game theory have seen significant strides in the past few years due to an increased interest in multi-agent reinforcement learning. Research using GA to inform policy decisions has continued recently in the context monetary economics [Arifovic et al. (2020)]. One notable example applying ABM (not GA) to policy-making was developed by Zheng et al. [Zheng et al. (2020)]. The work formulates the issue of optimal taxation as an OpenAI-Gym-style problem and uses deep reinforcement learning to model both decision makers and policy makers [Brockman et al. (2016)].

Taking inspiration from previous approaches to social learning and the recent advancements in reinforcement learning, this paper aims to outline a framework on which social learning can take place in tandem with individual learning. This comes as an effort to supplement the lack of a social aspect to learning in the AL literature while also allowing for a more detailed treatment of the learning process that happens in economic environments. In order to conform with the intuition that agent's rationality is bounded by the information they have been given, a fitting procedure is outlined using adaptive dynamic programming techniques that allow the study of learning behaviours when agents start with no knowledge of environment dynamics or rewards.

B. Learning Modelling

Ant Colony Optimization (ACO) is defined as a metaheurisitc general workflow that can be modified to solve several optimization problems. Drawing inspiration from the behaviour of wild ants, the workflow is characterized by autonomous agents that traverse the state space leaving behind pheromones. Pheromones can in turn influence the decision of future agents. Pheromones mix with known transition costs to guide agents to a decision; these pheromones aim to incorporate future costs, such that prospective ants avoid short term decision-making.

This class of swarm intelligence algorithm was first used to approximate a solution to the Traveling Salesman Problem (TSP), where the objective is to find the shortest path on which all nodes in a given graph are visited. ACO has previously been used in place of dynamic programming to approximate a solution to the problem, as it offers better scaling as state spaces grow. As such, it has become a common alternative to solve $\mathcal{NP}$-hard problems.

However, economic applications do not conform to several aspects of the vanilla TSP that motivated the original algorithm. The most prevalent difference lies in how the economic problems we are trying to study are stochastic processes; whereas state-transition costs are fixed in the TSP. Still, the metaheuristic has been used in the context of economics for portfolio optimization, option pricing and job shop scheduling [Hsu (2014)] [Deng and Lin (2010)] [Kumar et al. (2009)] [Huang and Liao (2008)]. This work also takes inspiration from earlier work done using ACO to find stochastic shortest paths and approximate solutions to the dynamic traveling salesman problem [Horoba and Sudholt (2010)] [Guntsch et al. (2001)]. Although unrelated to economics, these problems are comparable to dynamic inter-temporal decision-making in macroeconomic models.

The general ACO workflow is encompassed in Algorithm 1 followed by a discussion of formulas and parameters common to most uses of the metaheuristic. We outline modifications made to this framework in Section III.

Algorithm 1: Ant Colony Optimization Metaheuristic

```
initialization
while termination condition is not met do
    for each ant generated do
        construct solution
    for each path history and solution do
        update pheromones
    pheromone evaporation
```

This framework offers great flexibility in how certain steps are implemented, such that it can be modified to approach numerous problems[2].

In the original Ant System (AS) formulation by Dorigo Dorigo et al. (2006), heuristic information is weighted with pheromones to form transition probabilities, which can be used as stochastic weights in a decision making step. This weighting is shown is equation 5.

$$p(s,a) = \frac{\tau(s,a)^{\alpha} \cdot \eta(s,a)^{\beta}}{\sum_{a' \in A(s)} \tau(s,a')^{\alpha} \cdot \eta(s,a')^{\beta}} \tag{5}$$

where $A(s)$ is the actions available at state s and η are the known immediate rewards from taking an action at a given state.

The parameters α, β and ν are commonly preserved independent of the problem at hand; their original usage is shown in equation 5. The α parameter dictates the degree of importance assigned to pheromones; similarly, β controls the degree of importance assigned to known transition costs (broadly referred to as heuristic information). Finally, ν — called the evaporation rate — defines the rate on which past pheromones fade as time passes, making way for newer information. The last parameter is a key part of updating pheromones, which is done as follows:

$$\tau'(s,a) = (1-\nu) \cdot \tau(s,a) + \sum_{k=1}^{m} \Delta\tau(s,a) \tag{6}$$

where $\tau(s,a)$ is the current pheromone for state-action pair (s,a) and $\tau'(s,a)$ is its update; m is the number of ants on which the state-action pair applies to; and $\Delta\tau(s,a)$ is the pheromone left by ant k.

In the context of economic decision making and information dynamics, α and β can be interpreted as the influence decision-making agents give to socially-transmitted information versus long-running beliefs. Likewise, the evaporation rate ν can be seen as the degree on which socially-transmitted information becomes less prominent with time, or the rate on which this information evolves into long-standing beliefs.

[2]In the following equations, we slightly deviate from original notation in order to preserve compatibility with the problem discussed in this paper. Notably, node-to-edge subscript notation (τ_{ij}) to state-action notation ($\tau(s,a)$), where states are interpreted as nodes and actions as edges.

III. Model

A. Dyna-Q

While ACO can model social logic there must also be a mechanism to fit more dynamic models. One suitable choice is the Dyna-Q architecture, due to its compatibility with ACO and its potential economic intuition [Sutton (1991)]. A reasonable choice is a tabular environment with infinite or finite episodes[3].

Algorithm 2: Basic Dyna-Q Algorithm

Initialize $Q : S \times A \mapsto \mathbb{R}$, $M : S \times A \mapsto S \times \mathbb{R}$
while *Condition not met* **do**
  Sample trajectories $\tau_1, \tau_2, \dots \tau_n$ from the environment using policy π.
  Fit $M(S, A)$ using trajectories $\tau_1, \tau_2 \dots \tau_n$.
  for *Planning steps* **do**
    Sample S_{train}, A_{train} randomly from within $\tau_1 \dots \tau_n$
    $R_{train}, S'_{train} \sim M(S_{train}, A_{train})$
    $Q(S_{train}, A_{train}) \leftarrow$
     $Q(S_{train}, A_{train}) + \alpha \left(R_{train} + \gamma \max_A Q(S'_{train}, A) - Q(S_{train}, A_{train})\right)$

This paper implements the Dyna-Q+ algorithm in order to add a time-incentive to explore. This algorithm builds on what we have seen by adding a time-weight parameter κ that artificially increases rewards for older or unseen state-action pairs. The algorithm is the same as Algorithm 2 with two modifications: a) additional variables keep track of when a state-action pair was last seen and b) the following is added to R_{train}

$$R_{train} = R_{train} + \kappa \cdot \sqrt{\Delta t}$$

where Δt is the time since the last time this state-action pair was seen.

The economic intuition behind the choice for the Dyna algorithm lies in how new information is "pooled" into shared models of the economy, which is then used to inform agents of the market landscape. Additionally, the information is integrated into public perception based on trajectories that are not necessarily complete. A real world example of such a model is a newspaper, which surveys common knowledge and past experience to inform how economic agents go on to perceive value.

The Dyna-Q framework also allows us to separate social from individual learning unlike previous macroeconomic models. By keeping the model M and rewards Q separate, we essentially modularise the social and individual learning components of our model. While socially-shared experiences update the shared model of the economy, samples from this model can be used to train individual models, possibly with supplemented private data. Although Algorithm 2 and simulations in Section IV use Q-tables as individual models, it would be possible to use recursive least squares or other techniques commonly used in the AL literature.

B. The Algorithm

Our main algorithm is composed of iterations of generation, decision and updating, outlined in more detail in Algorithm 3. Generation is the process of generating new, finitely-lived agents.

[3]Using a different model-based approach such as MBPO Janner et al. (2019) could allow us to optimize on a continuous environment. The discussion and implementation of this method are beyond the scope of this paper, but likely a subject of future research.

Decisions are made by each agent based on Q values. Each agent will experience the environment differently and keep a history of the rewards earned in their lifetime. The joint history of all agents in a period will be fed to model M, which will in turn influence decisions made by future agents. Other mechanics are based on the Dyna-Q+ algorithm, as we can see in algorithm 3.

Algorithm 3: Main Information Transmission Routine

Initial $Q : S \times A \mapsto \mathbb{R}$, $M^{pher}, M^{belf} : S \times A \mapsto S \times \mathbb{R}$, Time Weight κ
while *Condition not met* **do**
  for *Number of ants* **do**
    Sample trajectories $\tau_1, \tau_2, \ldots \tau_n$ from the environment using policy π.
    Fit $M^{pher}(S, A)$ using trajectories $\tau_1, \tau_2 \ldots \tau_n$.
  for *Planning steps* **do**
    Sample S_{train}, A_{train} randomly from within $\tau_1 \ldots \tau_n$
    $R_{train}, S'_{train} \sim M^{join}(S_{train}, A_{train})$
    $Q(S_{train}, A_{train}) \leftarrow$
     $Q(S_{train}, A_{train}) + \alpha \left(R_{train} + \gamma \max_A Q(S'_{train}, A) - Q(S_{train}, A_{train})\right)$
  Evaporation

A general visualization of the model is presented in Figure 1:

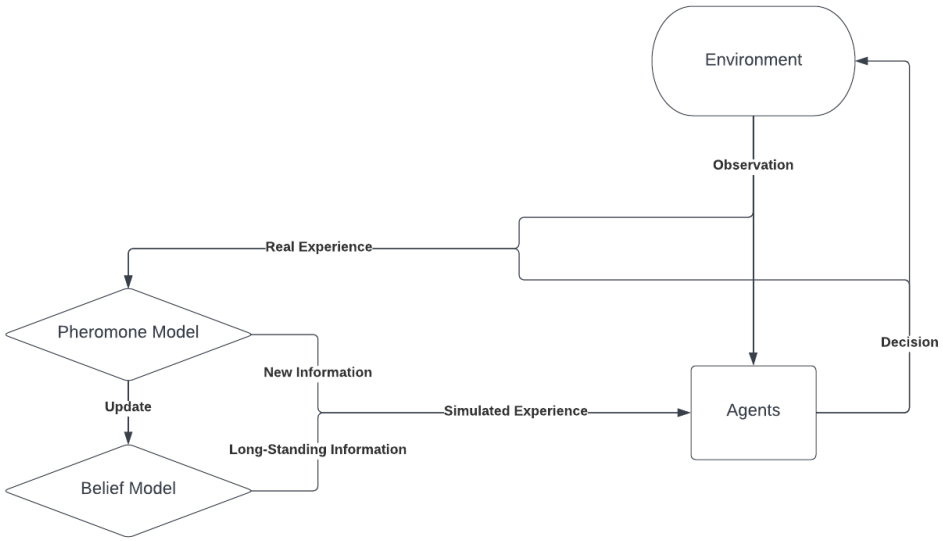

Figure 1. : Model Diagram. The figure shows the basic flow of information through the proposed model. Information begins as observations that are acted upon by an agent and used as training data for a pheromone model. Agent actions set a prior for upcoming observations from the environment. Pheromone models update belief models through evaporation and agents learn from simulated experiences sampled from both of them. Learning will in turn affect how agents make decisions in the future.

C. Generation

Autonomous agents are at the core of the mechanics outlined in this paper, as they allow for a study on the gradual change of beliefs in an environment. Agents need not be homogeneous, however, and can employ any mixture of parameters in order to model a scenario; as outlined in Section II.B, said parameters control agents' perception of the environment and influence they hold on it. Results from Section IV explore a fairly limited scope of parameters, which warrants future work exploring the use of stochastic agent generation and model-mixing in describing learning dynamics.

D. Sampling

Drawing from the ACO metaheuristic, the model in this work separates beliefs from pheromones. Beliefs are interpreted as long-standing expectations of heuristic information and used in conjunction with pheromones in order to form a decision. Model predictions are based on a combination of beliefs and pheromones. This combination is done for both rewards and transition dynamics, which together yield the sample from M^{join}. Here, we separate the reward and transition dynamics pheromones and beliefs, which we denote as R^{pher}, R^{belf}, D^{pher}, D^{belf} respectively. Notice that although we maintain the exponential weighting of dynamics presented in Section II.B, we use a weighted sum of rewards in order to maintain reward predictions similar to those seen in the environment. State-action rewards and dynamic predictions are defined by:

$$R_{train}(s,a) = \frac{\alpha \cdot R^{pher}(s,a) + \beta \cdot R^{belf}(s,a)}{\alpha + \beta} \tag{7}$$

$$p(S'_{train}|s,a) = \frac{D^{pher}(S'_{train}|s,a)^{\alpha} \cdot D^{belf}(S'_{train}|s,a)^{\beta}}{\sum_{s' \in S'(s)} D^{pher}(s'|s,a)^{\alpha} \cdot D^{belf}(s'|s,a)^{\beta}} \tag{8}$$

Models used in this paper employ an ε-greedy deterministic decision procedure, where actions with highest expected rewards are selected $\varepsilon\%$ of the time and randomly elsewhere. Decision are made as follows:

$$a_{t+1} = \begin{cases} \arg\max_a Q(s_t,a) & \epsilon < \varepsilon \\ P(a) = Q(s_t,a) & \text{elsewhere} \end{cases} \tag{9}$$

where $\epsilon \sim U(0,1)$

E. Updating

Updating is the process of changing beliefs on the expected state-action rewards in the current environment based on the agent's lifetime experience. This stage is split into two further parts: fitting and evaporation.

Fitting

The model fitting procedure is based on producing standard Monte-Carlo estimates for rewards and dynamics. One twist is that if we have not seen a certain state-action pair, its transition

probabilities are set to be uniform, as to avoid errors when combining pheromones with beliefs.

Evaporation

Again borrowing from the ACO metaheuristic, we use the concept of evaporation to bring pheromones closer to established beliefs. In addition to that, we use evaporation to allow established beliefs to incorporate newly-acquired information. The evaporation stage is defined by the following two operations:

$$R^{pher} = \begin{cases} R^{pher} & \text{if visited} \\ \kappa & \text{otherwise} \end{cases} \tag{10}$$

$$M^{belf} = (1-\nu) \cdot M^{belf} + \nu \cdot M^{pher} \tag{11}$$

This mechanism allows for general beliefs to converge to long-standing pheromone feedback, such that a "new normal" can be established after a certain amount of iterations following a change in dynamics.

IV. Simulations

A. Experimental Design

In order to test the proposed procedure, we implement environments to simulate the dynamics of known economic models. As the simulations proposed here are tied to certain economic models, the simulations' expressiveness might be limited by what that model aims to describe. That is, although environments used in ABM typically include specially-tailored dynamics not necessarily linked to standard economic models, we aim to remain faithful to our environments' original descriptions.

We propose the implementation of the McCall model of unemployment and the Huggett model of consumption following an OpenAI-Gym-style environment similar to the approach by Zheng et al. [Zheng et al. (2020)] [McCall (1970)][Huggett (1993)][Brockman et al. (2016)].

In the McCall model, agents receive a wage proposal at each step and can either accept or reject it. By default, accepting a wage terminates the episode and gives rewards equal to said wage up to the agent's age of death. The agent receives unemployment benefits if they are not employed in a certain episode, and dies at some pre-set age. This can be extended to let agents quit their jobs or have a chance of being fired at each episode. More formally, McCall agents face the following problem:

$$\max_{t'} \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t Y_t\right]$$
$$s.t.$$
$$Y_t = b,\ t < t' \text{ and } Y_t = w_{t'},\ t \geq t'$$

where β is the discount factor and Y_t is the income at time t. This income is equal to some constant unemployment benefit b or a wage $w_{t'}$ accepted at time time t'. The wages follow a beta-binomial distribution, that is, $w \sim BetaBin(n, \alpha, \beta)$.

Similarly, in the Huggett model, agents receive some stochastically-generated income at each step and can use it together with their existing assets to consume goods and gain utility from doing so[4]. Leftover assets or income are compounded by some interest rate and become the assets available in the next period. The agent can also borrow money up to a certain amount, so assets can also be negative. Interest is accrued on debt, which simply increases the amount of debt the agent has. The problem is formulated as the consumer problem:

$$\max_{c_t} \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t u(c_t)\right]$$
$$s.t.$$
$$a_{t+1} + c_t \leq Y_t + (1+r)a_t \text{ and } a_t \geq -B$$

where, c_t is consumption at time t, a_t are assets at time t, r is a constant interest rate, B is a constant borrowing constraint and $Y_t \sim BetaBin(n, \alpha, \beta)$ is an income shock at time t.

For each of the environments, we test four separate algorithms: Q-Learning, DynaQ, DynaQ+ and ACO DynaQ+.

We aim to compare how these perform in terms of speed, fitting patterns, optimal Q-values and average lifetime utility with greedy policies. It is important to mention that environment-specific intuitions shall be applied to some of these analyses. Firstly, optimal Q-values should reflect some sort of economic conclusion these models aim to answer. In the case of the McCall model, the moment where accepting a wage (action 0) brings higher expected utility than rejecting it (action 1) is called the *reservation wage* and should happen only once. In the case of the Huggett model, we should observe that higher asset levels warrant higher spending, and that lower asset levels risk bankruptcy, thus calling for reduced spending. Secondly, a sanity test for any of our algorithms in the Huggett model is whether an agent goes bankrupt, that is, gets -2000 utility at some point in the test. This has obvious signs in the results, and shouldn't happen in a well-fit algorithm. Note that this might still happen often during training, as bankruptcy is achievable at any state and not much exploration is needed to get there. The appendix contains a complete list of initialization parameters for each object used in the experiment.

B. Results

We begin by considering the comparaive speeds between the algorithms:

	Mean Rewards		Fitting Time	
Environment	Huggett	McCall	Huggett	McCall
Optimizer	(Utility)	(Income)		
Q-Learning	-1.698124	410.865750	84.621856	81.025059
DynaQ	-1.695528	417.330448	480.070568	470.455797
DynaQ+	-1.695875	422.447582	506.255441	494.747565
ACO DynaQ+	-1.756124	405.878000	92.554653	90.506384

Table 1—: Big Run Results. Values show the test-time performance of agents after $100,000$ Q-value updates.

[4]Here, we use a CRRA utility function with parameters described in the Appendix. A small quantity is also added to the denominator such that the punishment for zero consumption is non-infinite

	Mean Rewards		Fitting Time	
Environment Optimizer	Huggett (Utility)	McCall (Income)	Huggett	McCall
Q-Learning	-1.687222	410.362499	7.350350	7.597060
DynaQ	-1.698818	408.935622	44.519349	45.034055
DynaQ+	-1.689297	417.272493	47.351125	46.833290
ACO DynaQ+	-1.888333	402.394940	8.287741	8.021743

Table 2—: Medium Run Results. Values show the test-time performance of agents after $1,000$ Q-value updates.

	Mean Rewards		Fitting Time	
Environment Optimizer	Huggett (Utility)	McCall (Income)	Huggett	McCall
Q-Learning	-1.833195	379.658680	0.094567	0.076132
DynaQ	-1.696222	396.769244	0.460102	0.467019
DynaQ+	-1.720275	390.504100	0.487222	0.493065
ACO DynaQ+	-2.068486	392.217915	0.101478	0.100523

Table 3—: Small Run Results. Values show the test-time performance of agents after 100 Q-value updates.

The tables were gathered from running each algorithm with $100,000$ Q value updates, which we refer to as our "big run". We compare this run with two smaller runs, which we will refer to as our "medium run" ($1,000$ Q updates) and "small run" (100 Q updates). Note the ACO DynaQ+ algorithm has fewer iterations (we divide the number of iterations for the other algorithms by the "Ants" parameter) in order to equate it to other algorithms in terms of Q value updates. As can be seen, the Huggett model has consistently higher fitting times than the McCall model. This is to be expected, as the Huggett model is an infinite-horizon model, while the McCall model terminates when a wage is accepted. Additionally, we can clearly see how the model fitting and sampling come with an extra cost in terms of run speeds, although this cost seems to be decreased in the ACO DynaQ+ algorithm. This is clearly due to the reduced number of iterations used in the algorithm, although it does say something about the approach as an in-between for Q-Learning and Dyna algorithms. The last point is further backed by our performance results, which seems to favour the original Dyna Algorithms. This holds in the big run but favours the ACO Dyna variation when the run is short. Note that regardless Q-Learning still performs best in all runs.

The test performances give a positive picture of our algorithms, although it undermines the need for the Dyna-like algorithms. Our tests are composed of 100 runs of 100 episodes using a purely-greedy function, which always picks the action with highest Q value. Firstly, we notice that given enough time, our algorithms are able to perform better on the models, as seen by the increase in mean rewards over larger runs. Most importantly, all algorithms given a medium or larger number of iterations were able to learn how to avoid bankruptcy and how to better manage one's income. As we shall see, these come at odds with our training curves, although it might just be evidence of the overall volatility of the models and need for more "fragile" exploration. Finally, we note that Q-Learning seems to consistently perform at the same level as the original Dyna variants, with our

ACO DynaQ+ variant lagging behind. This is only different in the McCall model in the small run case, where our variant out-performs other methods.

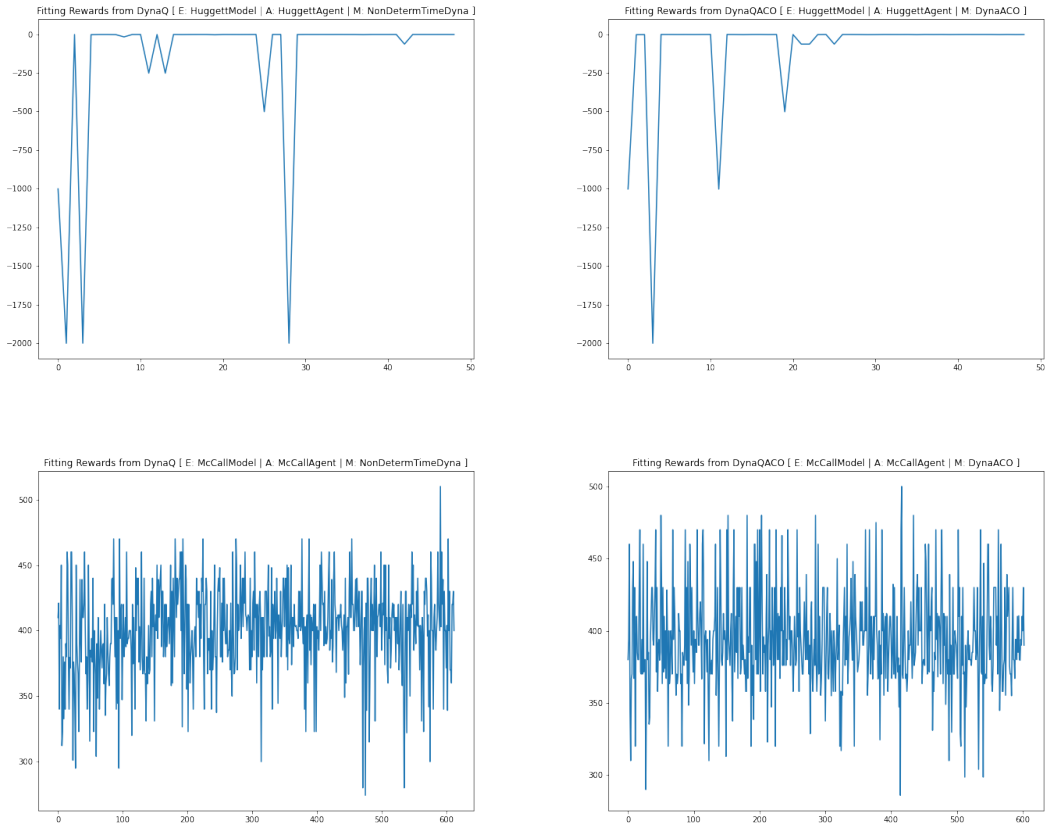

Figure 2. : Selected training curves for medium run. The graphs show the progression of agents' average utility per episode during training. The x-axis represents the episode number and the y-axis represents the average utility achieved in that episode (more is better). Downward spikes in the Huggett model show instances in which agents have no available income to spend in a given period and thus incur a penalty of −2000 utility in that period.

Overall, our fitting patterns are discouraging. As seen the following figures, our medium run fitting curves do not seem to be evidence of much learning. There does seem to be a slightly-upwards trend in some of the fitting curves, although even these do not show clear evidence of learning. Below we display the results of our DynaQ+ ACO model compared to the DynaQ+ algorithm in both environments after a medium run. The noise in the McCall model is to be expected, as state transitions are purely stochastic if a wage is rejected. As for the McCall model, exploration can cause severe dips in the model due to bankruptcy. This causes significant visual deviations in our fitting curves, although the main test should be on how agents act at test-time. Possible explanations aside, these results will put into question the effectiveness of using approaches based on Q-Learning when fitting economic models.

Our Q value interpretation seems to be more promising from the results in the big run, as seen in the figures. Although they are not consistent throughout the algorithms as we would expect from such a significant amount of time to learn the environment. Of the models, the Q-Learning

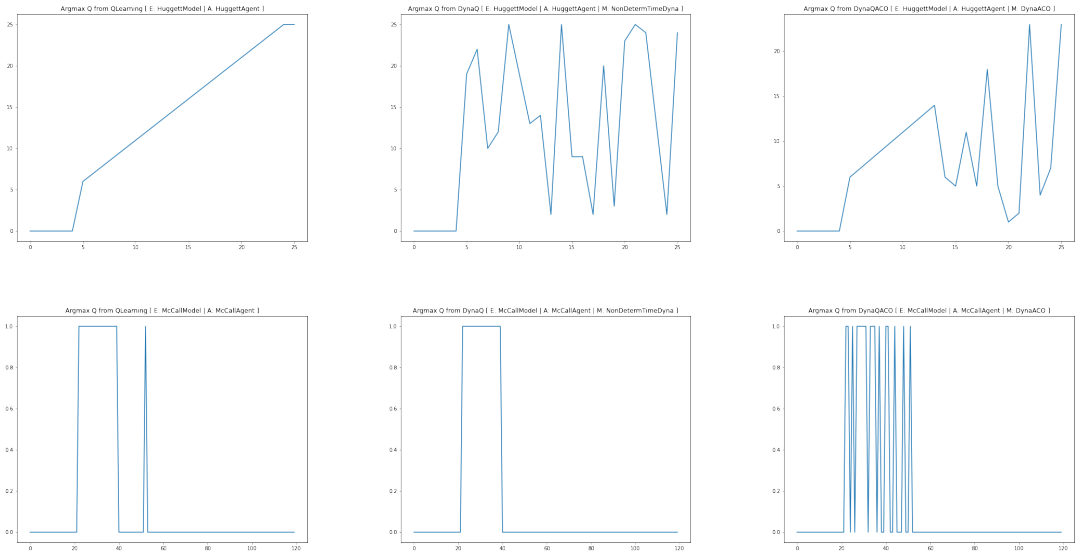

Figure 3. : Argmax over selected Q values following large run. The values represent the optimal actions at a given state as per the Q-table after extensive training. The x-axis shows a continuous state sequence and the y-axis shows the optimal learned action at that state. States in the Huggett model are asset levels and actions are the spending amount. States in the McCall model are wage offers and actions are binary for accepting or rejecting offers.

interpretations are the "cleanest", which suggests that model bias might come at odds with our aim to answer economic questions; thus, we include the Q-Learning results below as well. The McCall model seemed to be the most challenging for algorithms to learn, as is evident from the Q value interpretations. Most algorithms are able to detect the theoretical shift at a proposed wage of 40, although most also have "spikes" elsewhere (except for DynaQ+,). Similar results are present in the Huggett model Q values, although a general upward trend is present in all results; Q-Learning and the DynaQ+ ACO algorithms are able to learn the full curve, which is encouraging. One explanation for the results of our Dyna-like algorithms is model bias and overly-reinforcing a few Q-values rather than exploring. Still, even DynaQ+ did not fully learn the curve, which might be indicative of a low time weight.

V. Conclusions

In this paper, we presented a model of information transmission using model-based reinforcement learning and mechanics from Ant Colony Optimization. We set up the theory to create said model, formulated it as an algorithm, and tested it in the context of two custom economic environments. We have found that although our algorithm satisfies the necessary economic intuition, it does not manage to outperform alternatives in the tested environments. Although our results provide evidence that agents' test-time decision-making conforms to economic intuition, training data shows unsatisfactory increases in utility from learning. In spite of these results, the question remains as to how to most effectively model learning in environments where agents do not know the model dynamics.

In comparison to the standard approach taken by AL and social learning using GA, our algorithms are clearly at a disadvantage. However, further work into the application of the methods outlined here is warranted. Perhaps the models themselves present unrealistic or over-stylized characteristics

that make our approaches especially ineffective. The McCall model, is very close to being a bandit problem, as the next state given a salary rejection is completely random. Similarly, the Huggett model introduces significant stochasticity to our action space, which might likewise imbalance the ease of learning. Perhaps models with increased complexity are necessary to better utilize the algorithms proposed; Although it is possible that the landscape of macroeconomic problems are incompatible with the tools presented in this paper.

REFERENCES

[Arifovic 1994] Arifovic, Jasmina: Genetic algorithm learning and the cobweb model. In: *Journal of Economic Dynamics and Control* 18 (1994), Nr. 1, S. 3–28. – Special Issue on Computer Science and Economics. – ISSN 0165-1889

[Arifovic et al. 2020] Arifovic, Jasmina ; Grimaud, Alex ; Salle, Isabelle ; Vermandel, Gauthier: Social learning and monetary policy at the effective lower bound. (2020)

[Brockman et al. 2016] Brockman, Greg ; Cheung, Vicki ; Pettersson, Ludwig ; Schneider, Jonas ; Schulman, John ; Tang, Jie ; Zaremba, Wojciech: *OpenAI Gym*. 2016

[Deng and Lin 2010] Deng, Guang-Feng ; Lin, Woo-Tsong: Ant Colony Optimization for Markowitz Mean-Variance Portfolio Model. 6466 (2010), 12, S. 238–245. ISBN 978-3-642-17562-6

[Dorigo et al. 2006] Dorigo, Marco ; Birattari, Mauro ; Stutzle, Thomas: Ant colony optimization. In: *IEEE Computational Intelligence Magazine* 1 (2006), Nr. 4, S. 28–39

[Evans and McGough 2020] Evans, George W. ; McGough, Bruce: Adaptive Learning in Macroeconomics. (2020), 11

[Farmer and Foley 2009] Farmer, J. ; Foley, Duncan: The economy needs agent-based modelling. In: *Nature* 460 (2009), 08, Nr. 7256, S. 685686

[Fisher 1911] Fisher, Irving: *The Purchasing Power of Money, its Determination and Relation to Credit, Interest and Crises*. Indianapolis, IN, USA : The Online Library of Liberty, 1911

[Guntsch et al. 2001] Guntsch, Michael ; Middendorf, Martin ; Schmeck, Hartmut: An Ant Colony Optimization Approach to Dynamic TSP. (2001), S. 860867. ISBN 1558607749

[Horoba and Sudholt 2010] Horoba, Christian ; Sudholt, Dirk: Ant Colony Optimization for Stochastic Shortest Path Problems. (2010), S. 14651472. ISBN 9781450300728

[Hsu 2014] Hsu, Chih-Ming: An integrated portfolio optimisation procedure based on data envelopment analysis, artificial bee colony algorithm and genetic programming. In: *International Journal of Systems Science* 45 (2014), Nr. 12, S. 2645–2664

[Huang and Liao 2008] Huang, Kuo-Ling ; Liao, Ching-Jong: Ant colony optimization combined with taboo search for the job shop scheduling problem. In: *Computers Operations Research* 35 (2008), Nr. 4, S. 1030–1046. – ISSN 0305-0548

[Huggett 1993] Huggett, Mark: The risk-free rate in heterogeneous-agent incomplete-insurance economies. In: *Journal of Economic Dynamics and Control* 17 (1993), Nr. 5, S. 953–969. – ISSN 0165-1889

[Janner et al. 2019] Janner, Michael ; Fu, Justin ; Zhang, Marvin ; Levine, Sergey: When to trust your model: Model-based policy optimization. In: *arXiv preprint arXiv:1906.08253* (2019)

[Kumar et al. 2009] Kumar, Sameer ; Chadha, Gitika ; Thulasiram, Ruppa K. ; Thulasiraman, Parimala: Ant Colony Optimization to price exotic options. (2009), S. 2366–2373

[Lin et al. 2013] Lin, William T. ; Tsai, Shih-Chuan ; Lung, Pei-Yau: Investors' herd behavior: Rational or irrational? In: *Asia-Pacific Journal of Financial Studies* 42 (2013), Nr. 5, S. 755776

[Lucas 1981] Lucas, Robert E.: *Rational Expectations and Econometric Practice: Volume 1.* NED - New edition. University of Minnesota Press, 1981

[McCall 1970] McCall, J. J.: Economics of Information and Job Search. In: *The Quarterly Journal of Economics* 84 (1970), Nr. 1, S. 113–126

[Muth 1961] Muth, John F.: Rational Expectations and the Theory of Price Movements. In: *Econometrica* 29 (1961), Nr. 3, S. 315–335. – ISSN 00129682, 14680262

[Sutton 1991] Sutton, Richard S.: Dyna, an integrated architecture for learning, planning, and reacting. In: *ACM Sigart Bulletin* 2 (1991), Nr. 4, S. 160–163

[Tversky and Kahneman 1982] Tversky, Amos ; Kahneman, Daniel: Judgment under uncertainty: Heuristics and Biases. In: *Judgment under Uncertainty* (1982), S. 320

[Zheng et al. 2020] Zheng, Stephan ; Trott, Alexander ; Srinivasa, Sunil ; Naik, Nikhil ; Gruesbeck, Melvin ; Parkes, David C. ; Socher, Richard: The AI Economist: Improving Equality and Productivity with AI-Driven Tax Policies. (2020)

APPENDIX

A1. Environment Parameters

	McCall		Huggett
Max Wage	60	Max Assets	20
Unemployment Benefits	25	Max Debt	5
—	—	Interest Rate	0.01
—	—	CRRA Gamma	1.5
Discount Rate (γ)	0.9	—	0.5

A2. Optimizer Parameters

	Q-Learning	DynaQ	DynaQ+	ACO DynaQ+
Learning Rate (α)	0.1	0.1	0.1	0.1
Exploration Rate (ϵ)	0.3	0.1	0.1	0.1
Planning Steps	—	50	50	50
Ants	—	—	—	50

A3. Model Parameters

	Vanilla Dyna	Time Dyna	ACO Dyna
Time Weight (κ)	—	1×10^{-4}	1×10^{-4}
Pheromone Weight (α)	—	—	0.5
Belief Weight (β)	—	—	0.5
Evaporation Rate (ν)	—	—	0.1

Pareto is a Canadian, faculty-led, scholarly journal dedicated to the publication of undergraduate economics research. Founded at University of Toronto Mississauga, the journal aims to showcase the diversity and quality of Canadian undergraduate economic research.

Website: uecutm.ca/pareto
Instagram: @pareto.ej

2023 Officers

Mary Kazek, University of Toronto Mississauga, *President*
Nina Zi Wei Low, University of Toronto Mississauga, *VP Marketing*

2023 Faculty Supervisor

Nicholas Zammit, University of Toronto Mississauga, *Editor and Faculty Supervisor*

Authors and Articles Featured

Gregory Chung, *Effects of Cannabis Legalization on Tobacco and Alcohol Prevalence*
Ravdeep Sandal, *Skilled Immigration and its Impact on Ontario Unemployment Rates*
Melissa Siqueira, *Estimating the "Breadwinner Impact" on Household Division of Labour*
Jun Takahashi, *The Effects of Unconventional Monetary Policy on Firm Capital Structure*
Renato Zimmermann, *Modeling Social Learning Using Dyna-Q and Ant Colony Optimization*

Office

Innovation Complex Room KN3274
3359 Mississauga Road
Mississauga, Ontario Canada L5L 1C6

PARETO
Undergraduate Journal of New Economists
University of Toronto Mississauga