

KOF Dissertation Series

Survey Based Research in Economics – Essays
on Methodology, Economic Applications and Long
Term Processing of Economic Survey Data

Matthias Bannert

KOF Dissertation Series, No 25, 2016

DISS. ETH NO. 23697

**Survey Based Research in Economics - Essays on
Methodology, Economic Applications and Long Term
Processing of Economic Survey Data**

A thesis submitted to attain the degree of
DOCTOR OF SCIENCES of ETH ZURICH
(Dr. sc. ETH Zurich)

presented by

MATTHIAS BANNERT

Diploma in Economics, University of Constance

born on 16.04.1982

citizen of Germany

accepted on the recommendation of

Prof. Dr. Jan-Egbert Sturm, examiner

Prof. Dr. Ulf-Dietrich Reips, co-examiner

2016

I never learn anything talking.
I only learn things when I ask
questions.

Lou Holtz

Preface

This thesis was written during my time as a researcher in the Business Tendency Survey Section of the KOF Swiss Economic Institute (ETH Zurich). I am thankful to my doctoral supervisor Prof. Jan-Egbert Sturm for his ongoing support and trust in my work. Without his detailed feedback and support alongside the generous freedom and opportunity he provided, it would not have been possible to work interdisciplinary in between the fields of economics, psychology, statistics and information engineering. I am also highly thankful to my co-supervisor Prof. Ulf-Dietrich Reips who was always willing to share his profound knowledge and experience in methodology and internet-based research. Furthermore, Prof. Reips deserves credit for fostering fruitful exchange with researchers outside the field of economics, particularly with his Constance *iscience* group and the international *webdatanet* group.

Alongside both of my supervisors, I would like to thank my co-authors Andreas Dibiasi, Dirk Drechsel, David Iselin, Heiner Mikosch and Samad Sarferaz all of whom contributed to this dissertation far beyond their contributions to our joint work. I would also like to thank Klaus Abberger and Richard Etter for their valuable advice and constantly open doors and their effort to make every day at KOF enjoyable. I am indebted to my lecturers, my colleagues in KOF and outside of KOF, who conducted surveys, tirelessly discussed my ideas, tested my code or simply shared their knowledge. Thanks to Christoph Basten, Carole Berset, Michael Dantlgraber, Barbara Frank, Hans-Martin Gaudecker, Nadia Genova, Michael Graff, Tim Kuhlmann, Marius Ley, Christoph Moser, Nicola Jordan, Pauliina Sandqvist, Norbert Schneider, Michael Siegenthaler, Boriss Siliverstovs, Stefan Stieger, Anne Stücker, Jayson Swanson, Charles Till, Paul Wiederkehr and Benjamin Wohlwend. I would also like

thank Tolga Coban and Charles Clavadetscher of the KOF IT department for their support and open-mindedness.

Further I would like to thank Kathrin Bloch, Vera Degonda and Nicole Koch for their administrative help and their support in layout questions respectively. I also thank Ioan Gabriel Bucur, Florian Eckert and Enes Poyraz for valuable research assistance. During my time as a researcher at KOF, I have greatly benefited from open source software and community support. Even though the list of significant contributors is far too long to single out specific persons I would like to thank Martin Maechler and Brian D. Ripley for their incredible contributions to the R language for Statistical Computing that has been an indispensable companion of my thesis. I would also like to thank Dirk Eddelbuettel whose packages have been an important basis for my own contributions.

I would like to thank my dear friends for their support, their tolerance, their advice, their competitiveness, their ability to listen and their honesty. Thank you, Andrea, Eva, Lisa and Jörg for sharing the struggle – you played a major role in finally getting it done. Thank you Reto, Ste, Thomas, Moni, Meret, Christian, little Hannah, Reto and Nati, you made us feel at home while working and living in Switzerland. The marathon that writing this thesis has turned out to be, has asked for timeouts, patience and endurance alike. Thank you Marie, Raphi, Dende, Britta, Toni, Stefan, Hannes, Sabine, Lelle, Fred, Fabi, Pini, Mac, Juan, Phil, for providing the much needed breaks and inevitable balance. Research and software development made me put my head down, ignore my surroundings and not return calls. Thank you, long time companions Stephan, Oli, Mira, Andre, Tally, Toby, Alex, Julia and Eva for never quitting on me and providing the retreat you did.

Much love to my wonderful wife Eva, my parents Dagmar and Christian, Peter and Eva and my entire family in Augsburg, Bremen, Hamburg, Munich, the Lake Constance area and in the U.S.. You have always been the basis of my work - especially when I was not able to see the incredibly long list of supporters.

Matthias Bannert

Contents

| | |
|---|----------|
| Summary of the thesis | ix |
| Zusammenfassung der Dissertation | xiii |
| 1 Introduction | 1 |
| 1.1 Surveys in economic research | 2 |
| 1.2 Goals, contributions and outline | 2 |
| 2 No Answer Is an Answer - What Drives Unit Non-Response in Business Tendency Surveys? | 9 |
| 2.1 Introduction | 10 |
| 2.2 Literature | 11 |
| 2.2.1 Non-response and missing data | 11 |
| 2.2.2 Meta-information on BTS | 13 |
| 2.3 Data | 14 |
| 2.3.1 KOF BTS in the service sector | 14 |
| 2.3.2 A survey about surveys | 15 |
| 2.3.3 Merged sample | 15 |
| 2.3.4 Variables | 17 |
| 2.4 Methodology | 21 |
| 2.4.1 Accounting for selectivity | 21 |
| 2.4.2 Estimation of unit non-response in a GLM framework | 22 |
| 2.4.3 Estimation of unit non-response in a two part model | 23 |
| 2.5 Discussion of estimation results | 24 |
| 2.5.1 Two part model estimation results | 24 |

| | | |
|----------|---|-----------|
| 2.5.2 | Variation and robustness of results | 28 |
| 2.6 | Conclusion and outlook | 28 |
| 3 | Macro and Micro Level Impulse Responses: A Survey Experimental Identification Procedure | 33 |
| 3.1 | Introduction | 34 |
| 3.2 | Method | 36 |
| 3.2.1 | Generating impulse responses from survey experiments | 37 |
| 3.2.2 | Firm level vs. aggregate impulse responses | 39 |
| 3.2.3 | Validity of the survey experiment | 43 |
| 3.3 | An application to oil price shocks | 44 |
| 3.3.1 | Survey set-up | 45 |
| 3.3.2 | Experimental design | 46 |
| 3.4 | Empirical results | 48 |
| 3.4.1 | Aggregate results | 48 |
| 3.4.2 | Heterogeneity within sectors | 51 |
| 3.4.3 | Industry level results | 53 |
| 3.4.4 | Firm-level results | 58 |
| 3.5 | Conclusion | 63 |
| 4 | How Are Firms Affected by Exchange Rate Shocks? Evidence From Survey Based Impulse Responses | 65 |
| 4.1 | Introduction | 66 |
| 4.2 | The Survey-Based Impulse Response Approach | 71 |
| 4.2.1 | Survey set-up | 72 |
| 4.2.2 | Experimental design | 73 |
| 4.2.3 | From firm-level to aggregate impulse responses | 75 |
| 4.3 | Results | 76 |
| 4.3.1 | Exchange rate expectations | 77 |
| 4.3.2 | Sector and Industry Level Effects | 77 |
| 4.3.3 | Firm-level heterogeneity | 86 |
| 4.3.4 | Firm-level regression analysis | 86 |

| | | |
|----------|--|------------|
| 4.4 | Conclusion | 91 |
| 5 | Ask Your Doctor or Pharmacist! On the Impact of Physician Dispensing on Utilization of Pharmacies | 95 |
| 5.1 | Introduction | 96 |
| 5.2 | Physician dispensing | 97 |
| 5.2.1 | General overview | 97 |
| 5.2.2 | State of legislation in Switzerland | 100 |
| 5.2.3 | Pharmaceuticals market in Switzerland | 102 |
| 5.3 | Data | 104 |
| 5.3.1 | Variables | 105 |
| 5.3.2 | Sample | 106 |
| 5.4 | Empirical strategy | 109 |
| 5.4.1 | Utilization of pharmacies | 110 |
| 5.4.2 | Measuring coverage | 110 |
| 5.5 | Results | 112 |
| 5.5.1 | Estimation results | 113 |
| 5.5.2 | Regional coverage patterns | 115 |
| 5.6 | Conclusion | 116 |
| 6 | Technical Aspects: A Software Package to Manage And Archive Economic Time Series Data | 119 |
| 6.1 | Introduction | 120 |
| 6.2 | Motivation | 122 |
| 6.3 | Data storage | 124 |
| 6.3.1 | Data: time series records | 125 |
| 6.3.2 | Meta data: localized and unlocalized information | 126 |
| 6.3.3 | Storing sets of time series | 126 |
| 6.3.4 | Mapping SQL relations to R objects | 127 |
| 6.3.5 | Function overview | 127 |
| 6.4 | Code examples | 129 |
| 6.4.1 | Write time series to database | 131 |

| | | |
|-------------------|---|------------|
| 6.4.2 | Read time series from the database | 133 |
| 6.4.3 | Adding and reading elaborate meta information | 133 |
| 6.4.4 | Plot a list of time series | 137 |
| 6.4.5 | Data explorer | 138 |
| 6.5 | Conclusion and outlook | 140 |
| Appendix A | | |
| | Appendix to chapter 2 | 143 |
| Appendix B | | |
| | Appendix to chapter 3 | 147 |
| Appendix C | | |
| | Appendix to chapter 4 | 167 |
| Appendix D | | |
| | Appendix to chapter 6 | 173 |
| | References | 176 |
| | Curriculum Vitae | |

List of Figures

| | | |
|-----|---|-----|
| 1.1 | Volatile participant with no remarkable pattern | 3 |
| 1.2 | Cruising participant with extremely low variance | 3 |
| 1.3 | Rookie participant with low within question variance | 4 |
| 1.4 | Dendrogram of response clusters | 4 |
| 2.1 | Unit Non-Response Count by Samples | 17 |
| 3.1 | Price of crude oil in CHF per barrel | 45 |
| 3.2 | Changes in real sales, foreign & domestic prices within 18 months . . . | 52 |
| 4.1 | CHF per EUR exchange rate | 67 |
| 4.2 | Exchange rate expectations within 18 months | 78 |
| 4.3 | Turnover 18 months vs. export shares | 80 |
| 4.4 | Costs 18 months vs. import shares | 81 |
| 4.5 | Turnover distributions | 87 |
| 4.6 | Costs distributions | 88 |
| 5.1 | Legislation by Canton | 100 |
| 5.2 | Dispensing physicians and pharmacies by Canton | |
| 5.3 | Locations of RoKA study participants | 107 |
| 5.4 | Pharmacies per 1000 citizen by regional typology (2012) | 108 |
| 5.5 | Density of dependent variables | 109 |
| 5.6 | Comparison of regional coverage in city centers of Fribourg and Lucerne | 116 |
| 6.1 | Overview: relational structure | 125 |
| 6.2 | R Studio password prompt | 130 |

| | | |
|------|--|-----|
| 6.3 | PostgreSQL client: a time series record | 132 |
| 6.4 | PostgreSQL client: a meta information record | 132 |
| 6.5 | Time series with dynamic meta description | 137 |
| 6.6 | Plot method for list of time series | 138 |
| 6.7 | Web based data explorer: build queries | 139 |
| 6.8 | Web based data explorer: plot and export | 140 |
| 6.9 | Distributions by samples | 152 |
| 6.10 | Aggregation scheme | 161 |
| 6.11 | Impulse response world-wide | 164 |
| 6.12 | Impulse response Swiss level | 165 |
| 6.13 | Responses of Swiss manufacturing real turnover | 166 |

List of Tables

| | | |
|-----|---|-----|
| 2.1 | Variable description | 18 |
| 2.2 | Hurdle part | 25 |
| 2.3 | Automatically Selected Variables vs. Initial Model | 29 |
| 3.1 | Survey based and BVAR impulse response functions | 50 |
| 3.2 | Signs of impulse responses to an oil supply shock | 55 |
| 3.3 | Industry level responses | 57 |
| 3.4 | Response of real sales, domestic and foreign sales prices to oil price shock | 61 |
| 4.1 | Industry and sector level survey based impulse responses | 84 |
| 4.2 | Changes in turnover, costs, and profits given exchange rate shock | 90 |
| 5.1 | Pharmaceuticals Market in Switzerland, source: interpharma.ch | 102 |
| 5.2 | Variable description | 106 |
| 5.3 | Utilization of Swiss Pharmacies | 114 |
| 6.1 | Student t-test on firms' projection errors | 150 |
| 6.2 | Mean oil shares of early and late respondents | 152 |
| 6.3 | Kolomogorov-Smirnov test results for equality of distributions | 153 |
| 6.4 | Variables (KOF Survey and external sources) | 159 |
| 6.5 | Mean export shares of early and late respondents | 169 |
| 6.6 | Kolomogorov-Smirnov test results for equality of distributions | 169 |
| 6.7 | Benchmark: Reading 100 time series, 100 replications | 175 |

Summary of the thesis

This thesis discusses surveys as a method to generate data for economic research. Chapters 2 and 6 focus on methodological and technological aspects of survey-based research while the 3rd, 4th and 5th chapter provide innovative application examples and address economic research questions. The introductory chapter is devoted to the general role of surveys in economic research and outlines the contributions of the thesis in that context.

Chapter 2 which is joint work with Andreas Dibiasi studies participant level drivers of unit non-response in business tendency surveys. We match standard business tendency survey data with firm characteristics, paradata and meta information obtained from an exclusive meta survey conducted by the KOF. We find that perceived response burden, a company's unit non-response history, participation mode and return time policy are the most robust drivers of unit non-response in our sample. Employing a two-part hurdle model provides a further explanation: We find that higher capacity utilization reduces the propensity to not return all four questionnaires within the observed timespan. Remarkably, we do not find a significant influence either of company size on the propensity to respond to periodical qualitative BTS or of participants' valuation of BTS.

Chapter 3, co-authored by Dirk Drechsel, Heiner Mikosch and Samad Sarferaz, analyzes the effects of macroeconomic shocks on prices and output at different levels of aggregation using a bottom up approach. We show how to generate firm-level impulse responses by incorporating experimental settings into surveys and by exposing firm executives to treatment scenarios. Subsequently, aggregation results in

industry-level and economy-wide impulse responses. We further show how the effects measured through survey experiments can be compared to impulse responses retrieved from VARs. We apply the procedure to study the effects of an oil price shock using a representative sample of over 1'000 Swiss firms. At the aggregate and industry-level our findings confirm, with some notable exceptions, results from a standard VAR. At the micro-level we analyze the driving forces behind firm-specific impulse responses, controlling for several firm characteristics via panel data analysis and thereby solving several existing puzzles.

Chapter 4, also jointly written with Dirk Drechsel, Heiner Mikosch and Samad Sarferaz, relies on the survey based framework developed in chapter 3. Based on a survey experiment, we study firm reactions to an exchange rate shock in the light of the recent 2012 and 2015 interventions of the Swiss National Bank (SNB). Following chapter 3, we use a survey-based approach to create impulse responses on various levels of granularity from executives' answers in response to a shock scenario. Based on our own scenario survey, which was conducted in 2012 while the exchange rate floor was still in effect, we find that sectors react heterogeneously: Manufacturing turnover would decrease by 3.3% within six months and by 4.3% within 18 months. Total costs would decline by 1.3% within six months and 2.0% within 18 months, while profits would shrink by 3.3% within six months. The effects are substantially lower for the service and construction sector, but exhibit large variation across sub-sector industries. Panel regression analysis reveals that firm-specific export shares and intermediate goods import shares are key determinants of firms' turnover, costs and profits reactions. Further the degree to which firms are shocked, measured by their unconditional exchange rate expectations influences firms' impulse responses.

Chapter 5 which is joint work with David Iselin provides another application of economic surveys and shows how surveys can be used to study specific questions. Chapter 5 intends to supplement the discussion on physician dispensing with another perspective: We match pharmacy-level survey data with data from several sources to study the relation between pharmacies' utilization and dispensing physicians. We exploit the fact that dispensation regulation is organized at the state- (canton) level in Switzerland to compare pharmacies that operate under different regimes

within a relatively small area. We find a substantial *ceteris paribus* reduction in the utilization of pharmacies in cantons that allow physicians to dispense drugs. Further, we find that pharmacies seem to adapt their role and portfolio in regions that allow physicians to dispense medication.

The final chapter discusses important technical aspects of conducting long term business tendency surveys. Chapter 6 suggests how data from periodical surveys could be managed in an efficient and transparent way that ensures availability and consistency over time while being maintainable for economic institutions. Chapter 6 introduces the *R* software package *timeseriesdb* which was developed over the course of this dissertation: *timeseriesdb* provides an interface between the popular relational database management system PostgreSQL and the R Language for Statistical Computing. The package suggests a setup that was designed to manage time series data stemming from surveys and establishment statistics. Hence, the package is optimized to handle a large amount of regular, rather low frequency (daily, monthly, quarterly or annual) time series as opposed to managing a smaller number of high frequency time series such as the information returned from measuring devices. The *timeseriesdb* package also provides the opportunity to store comprehensive multi-lingual meta information. The underlying structure solely uses license cost free open source components: The *R Language for Statistical Computing* and the advanced relational database management system *PostgreSQL*.¹ Chapter 6 of the dissertation describes the functionality of the package, offers applied examples and installation notes.

¹The *R* package *timeseriesdb* is freely available from the Official R package repository (CRAN): <http://cran.r-project.org>.

Zusammenfassung der Dissertation

Die vorliegende Dissertation beschäftigt sich mit Umfragen als Methode zur Gewinnung ökonomischer Daten. Dabei geht die Arbeit im zweiten und letzten Kapitel auf methodische bzw. technische Aspekte ein und zeigt in drei weiteren Kapiteln innovative Anwendungen von Umfragen zur Beantwortung ökonomischer Fragestellungen. Das Einführungskapitel diskutiert die Rolle von Umfragen in der empirischen ökonomischen Forschung und skizziert die Beiträge der vorliegenden Arbeit in diesem Kontext.

Kapitel 2, das in Zusammenarbeit mit Andreas Dibiasi entstanden ist, beschäftigt sich mit den Determinanten von Non-Response in Konjunkturumfragen auf Teilnehmerebene. Wir haben Konjunkturumfragedaten mit Firmencharakteristiken sowie Paradata und Meta Informationen, die aus einer eigens durchgeführten Feedbackumfrage gewonnen werden können zusammengeführt. Es zeigt sich, dass die wahrgenommene Umfragebelastung, die historische Non-Response Bilanz einer Firma, der Teilnahmemodus und die Firmenpolicy bezüglich der Antwortdauer die wichtigsten erklärenden Variablen für Non-Response sind. Ausserdem reduziert eine höhere Kapazitätsauslastung die Wahrscheinlichkeit, dass eine Firma nicht an allen vier Quartalsumfragen teilnimmt. Interessanterweise finden wir keinen signifikanten Einfluss der Unternehmensgrösse auf die Wahrscheinlichkeit an einer regelmässigen qualitativen Konjunkturumfrage teilzunehmen. Die Bewertung von Konjunkturumfragen seitens der Unternehmen hat ebenfalls keinen signifikanten Einfluss auf das Nicht-Antwort verhalten.

Kapitel 3, gemeinsam mit Dirk Drechsel, Heiner Mikosch und Samad Sarferaz geschrieben, analysiert die Effekte makroökonomischer Schocks auf Preise und Output auf unterschiedlichen Aggregationsstufen in dem das Kapitel einen Bottom-Up Ansatz wählt. Wir zeigen wie Impuls-Antworten auf Firmenebene gebildet werden können in dem man Experimentalanordnungen in Umfragen implementiert und Führungskräfte mit Szenarien konfrontiert. Durch Aggregation können anschliessend Resultate auf Industrie- oder Sektorebene berechnet werden. Wir zeigen ausserdem, wie die Resultate aus den Umfrageexperimenten mit VAR Resultaten verglichen werden können. Wir wenden das Verfahren auf ein Sample von mehr als 1000 Schweizer Firmen an, um die Effekte eines Ölpreisschocks zu studieren. Im Gesamtaggregat bzw. auf Industriebene bestätigen unsere Ergebnisse, mit einigen Ausnahmen, die Resultate von Standard VARs. Auf Mikroebene können wir mittels Paneldatenanalyse für Firmencharakteristika kontrollieren und die treibenden Kräfte hinter firmenspezifischen Impuls-Antworten analysieren.

Kapitel 4, ebenfalls mit den Co-Autoren des dritten Kapitels geschrieben, beruht auf dem in Kapitel 3 entwickelten umfragebasierten Framework. Basierend auf einem Umfrageexperiment untersuchen wir die Reaktionen von Firmen auf einen Wechselkursschock im Licht der jüngsten massiven Interventionen der schweizerischen Nationalbank (SNB) in den Jahren 2012 und 2015. Entsprechend der in Kapitel 3 entwickelten empirischen Strategie verwenden wir einen Umfrageansatz um Impuls-Antworten auf unterschiedlichen Aggregationsstufen aus den Antworten von Führungskräften auf ein Schock-Szenario zu generieren. Basierend auf unserer eigenen Szenariobefragung, die im Jahr 2012 durchgeführt wurde, als die Wechselkursuntergrenze noch bindend war, zeigt sich dass die einzelnen Sektoren heterogen reagieren: Umsätze im verarbeitenden Gewerbe sinken um 3.3% innerhalb von sechs Monaten und um 4.3% innerhalb von 18 Monaten. Die Totalen Kosten sinken um 1.3% innerhalb von sechs Monaten und um 2.0% innerhalb von 18 Monaten, während die Gewinne um 3.3% schrumpfen. Die Effekte im Dienstleistungs- und Bausektor sind deutlich geringer aber weisen massive Unterschiede innerhalb der Sektoren auf. Panelregressionanalyse zeigt dass firmenspezifische Exportanteile und Importe von Zwischenprodukten die entscheidenden Determinanten der Umsatz-,

Kosten und Gewinnreaktionen sind. Ausserdem beeinflusst der Grad zu welchem Firmen geschockt sind, gemessend anhand der ursprünglichen Wechselkurserwartungen, diese Kennzahlen.

Das fünfte Kapitel entstand gemeinsam mit David Iselin und ist ein Anwendungsbeispiel, das zeigt wie Umfragen genutzt werden können, um sehr spezifische Fragestellungen zu erforschen. Kapitel 5 versucht die Diskussion um direkte Medikamentenabgabe durch Ärzte um eine weitere Perspektive zu ergänzen: Wir führen Daten auf Apothekenebene mit Daten anderer Quellen zusammen, um das Verhältnis zwischen der Nutzung von Apotheken und Regelungen, die Ärzten erlauben Medikamente direkt abzugeben, zu untersuchen. Dabei nutzen wir die Tatsache, dass die Schweiz deartige Gesetze auf Kantonsebene implementiert und uns so die Möglichkeit gibt Apotheken, die in unterschiedlichen gesetzlichen Ausgangssituationen operieren innerhalb einer relativ kleinen Fläche zu vergleichen. Wir finden eine substantielle *ceteris paribus* Reduktion der Nutzung von Apotheken in Kantonen die Ärzten erlauben Medikamente direkt abzugeben. Ausserdem zeigt sich, dass Apotheken ihre Rolle und ihr Portfolio an die gesetzliche Lage anzupassen scheinen.

Das sechste und letzte Kapitel geht der Frage nach, wie die aus regelmässigen Umfragen gewonnen Daten nachvollziehbar und effizient im Bezug auf Verfügbarkeit, Verwaltbarkeit und Konsistenz als Zeitreihen gespeichert werden können. Kapitel 6 stellt das im Rahmen der Dissertation entwickelte *R Software* Paket *imeseriesdb* vor, das eine Schnittstelle zwischen der populären, relationalen Open Source Datenbank PostgreSQL und der *R Language for Statistical Computing* schafft.² Zudem wird eine Datenbankstruktur vorgeschlagen, die sich insbesondere auf die Speicherung mehrsprachiger Metainformationen im Kontext von Speicherung Zeitreihen aus amtlichen Statistiken und ökonomischen Umfragen eignet. Im Gegensatz zu Zeitreihendatenbanken die hochfrequente Daten, etwa von Messgeräten, aufnehmen, fokussiert sich *timeseriesdb* auf monatliche, quartalsweise und jährliche Frequenz und sowie kontextbezogene Beschreibung dieser Daten. Kapitel 6 der Dissertation

²Das *R* Paket *timeseriesdb* is frei verfügbar und kann aus dem offiziellen *R* repository <http://cran.r-project.org> heruntergeladen werden.

beschreibt die Funktionsweise und bietet praktische Anwendungsbeispiele zur Verwendung der Software sowie eine Installationsanleitung.

Chapter 1

Introduction

1.1 Surveys in economic research

To pose questions seems to be an apparent and literal way of conducting research – particularly in the social sciences. This also holds for empirical economic research: Though lab experiments and non-reactive methods have become popular, surveys continue to play an important role in empirical economics. Despite the massive emergence of alternative data sources and opportunities to track economic agents, surveys remain a sound and reliable method to gather timely information. Information obtained from economic surveys is flexibly used in a wide array of applications: academic research, economic forecasts and establishment statistics. Particularly business tendency surveys (BTS) have a long tradition and are well established in many countries, as they help to monitor a country’s business cycle.

The focal role of survey data implies the need of a high-quality data generating process. Moreover forecasting and many other econometric exercises such as vector auto regressions (VAR) demand longitudinal data. In turn the responsibility of survey conductors to ensure data quality is not a static task. Captivating advances in information technology as well as social development have been influencing research methods, participant behavior and data management alike. In order to keep up with these developments economic researchers need to adapt survey methodology and monitor participant behavior over time while accounting for the idiosyncratic nature of economic surveys. Economic survey researchers are facing long-known and new challenges, yet at the same time modern surveys also provide new opportunities. Though this thesis intends to address challenges of modern survey-based research, it is also motivated by the idea of illustrating strengths and opportunities of surveys as a research method in modern empirical economics.

1.2 Goals, contributions and outline

The aim of this dissertation is twofold. First, this thesis intends to provide a multi-faceted view on surveys in economics – from participant behavior (chapter 2) to innovative opportunities of the method to the technical aspects of longterm data management. Second, the applications presented in chapters 3 to 5 contribute to

the existing economic literature: chapter 3 discusses firms' reactions to an oil price shock in an impulse response setup, chapter 4 contributes to the recent exchange rate debate in Switzerland and chapter 5 discusses a health economics related topic. Besides these primary goals, the sixth and final chapter of the thesis intends to add a technological contribution.

The following chapter, written together with Andreas Dibiasi, intends to update the business tendency survey specific look at non-response. Further, we aim at a participant-level analysis of data quality and response behavior as opposed to survey wave-level metrics of data quality. If we picture qualitative response at the participant-level we can literally think of a pattern. Figure 1.1 , through and figure 1.3 show exemplary visualizations of participant-level responses.

Figure 1.1: Volatile participant with no remarkable pattern



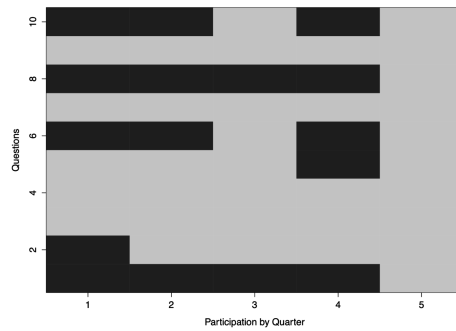


Figure 1.2: Cruising participant with extremely low variance



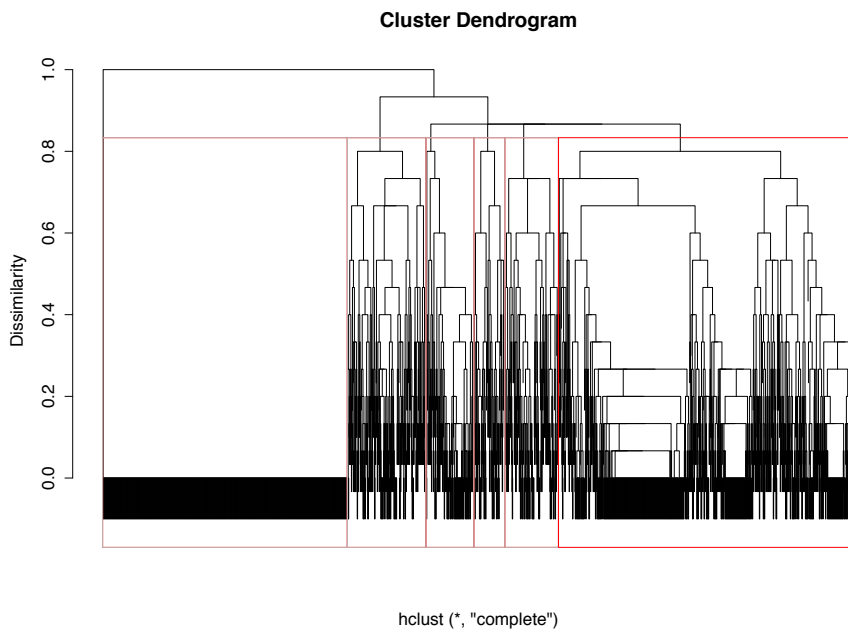
The X-axis denotes time, sequentially counting the quarters of participation and the Y-axis depicts question indices. Light areas indicate positive, medium areas

Figure 1.3: Rookie participant with low within question variance



depict neutral and dark areas depict negative answers. White areas indicate item non-response. These figures show the prevalence of response patterns in an intuitive fashion, each graph represents one participant over time and all questions. Figure 1.4 visualizes a computing driven effort to group patterns of survey responses with a cluster analysis. The large cluster on the left side is the tightest and most similar cluster of answers and represents unit non-response.

Figure 1.4: Dendrogram of response clusters



Chapter 2 focuses entirely on this non-response cluster. Despite a large existing strand of literature on non-response, it is easy to motivate further research for two main reasons. First, the way humans feel committed to official requests in general and particularly in business may have changed dramatically over time. Thus for survey conductors courting the attention of their participants, it is important to study and monitor the development of response behavior on a regular basis. Second, while most of the aforementioned literature stems from the fields of psychology and survey methodology, contributions from economists are rather scarce. Nevertheless, it is important to address economic surveys specifically because of their idiosyncratic nature. Dillman (2000) suggests that response behavior in business surveys needs to be studied specifically because participants answer on behalf of their company, as opposed to answering questions addressed at an individual. Also, unlike most individual surveys, business tendency surveys continue to ask the same mostly qualitative questions on a regular basis. Thus results from other fields of survey-based research cannot be carelessly transferred to economic surveys. Chapter 2 contributes to methodological literature not only by adding an economic perspective but also by addressing the participant level. Due to this focus we also collected additional data by conducting a feedback survey that was explicitly designed to query additional firm-level characteristics that have the potential to explain non-response at the participant level. We study how variables such as a company's size, business situation, non-response history, response burden or valuation of surveys influence non-response and with it the reliability of economic surveys.

Chapter 3, which is based on a paper with Dirk Drechsel, Heiner Mikosch and Samad Sarferaz introduces a survey-based approach to generate firm level impulse response to study the effects of an oil price shock. The goal of chapter 3 is to supplement existing state of the art VAR-based approaches that only allow insights at the macro-level. Our approach allows generating impulse responses at more finely grained levels. We chose a well-studied variable to show that our survey-based impulse responses are consistent with robust state of the art methods when aggregated. A questionnaire can be a rather intuitive way of extracting causal reactions when other methods live on assumptions and identifying restrictions. In a survey, it is

comparatively easy to set up counterfactual scenarios and query expectations and reactions under different regimes from the same participant. Unlike in an observational study, a scenario survey generates data for two mutually exclusive states and can thus help to evaluate the impact of a regime change. Additionally, surveys are not only an alternative but can be regarded as as a complementary method that may help to test identifying restrictions. Chapter 4 continues to apply the same method to address a different economic question. In light of the most recent interventions of the Swiss National Bank (SNB) in 2012 and 2015, we apply the survey-based impulse response approach to assess the effect of an exchange rate shock in the Euro / Swiss franc exchange rate. We contribute to the lively public and academic debate about the effects of a substantial exchange rate shock on the Swiss economy by providing results at different aggregation levels. Using firm level survey based impulse responses we are able to show that reactions are very heterogenous across sectors. Chapters 3 and 4 are based on an explicitly designed ad hoc survey that was pre-tested with selected participants and afterwards attached to the KOF investment survey.

Chapter 5 intends to contribute to the political debate on drug dispensation regulations in Switzerland and provides an example for the flexible use of survey data to address economic questions. The fifth chapter of the thesis, co-authored by David Iselin, combines an online survey among Swiss pharmacists, which was originated as a business study, with municipality level statistics and data from Google's Geocoding and Places application programming interfaces (APIs). By combining survey data with data from a non-reactive data generating process, chapter 5 provides an example for the supplementary use of survey data with data stemming from other data generating techniques. By studying an unique dataset we supplement the traditional health expenditure-focused debate, which is mostly based on insurance data, with another perspective. We provide pharmacy-level information and show how utilization of pharmacies is related to legal regulations.

In addition to its methodological and economic research questions, this dissertation has always worked on the technological aspects of survey-based research. This interest in the technology of a survey's ecosystem is motivated by two aspects:

First, in the sense of Buckheit and Donoho (1995) this thesis appreciates scientific computing and software development to be a part of conducting, documenting, communicating and replicating data driven research. Second, ensuring data quality in long-term research projects requires researchers to understand the data generating process as well as the processing of data thoroughly to guarantee transfer of knowledge among researchers. Consequently, economists and statisticians Koenker and Zeileis (2009) argue that the development of scientific software can no longer be explicitly left to professional software vendors.

Given this understanding of tech savvy researchers, the final chapter intends to encourage economists to incorporate the concept of what has been summarized as *Reproducible Research* in the literature. Yet, the main contribution of the chapter is the software package it introduces. Chapter 6 introduces an open source software package to transparently manage large amount of time series data aiming at time series generated by economic surveys and establishment statistics in particular. The long tradition of survey based research and the demand for consistent longitudinal data confronts survey researchers with the challenging task of long-term data management. In many research contexts, due to economic time series' long life time, data generating processes, as well as processing of the resulting data itself may consist of a hodgepodge mix of software and ancient processes that hinder reproducibility. Furthermore, for privacy reasons, survey data are often only exchanged at aggregated levels, making researchers at the receiving end dependent on context-aware meta information. The software package introduced in chapter 6 provides the economic researcher with the opportunity to get involved in data management at an early stage more easily, and to manage data and meta information in a single consistent, ecosystem.

The package was developed over the course of this dissertation and consists entirely of license costs free open source components: a schema for the relational database management system (DBMS) *PostgreSQL* and the *R* package *timeseriesdb*. By choosing a light-weight setup that is easy to maintain, the approach suggested in chapter 6 is also suitable for smaller institutions and chairs which are seeking to

manage their data centrally. The choice of the *R Language for Statistical Computing* as the main processing language for the time series data themselves is not only justified by language's growth rates and popularity level that are unparalleled by other domain-specific languages, but mostly due to the fact that *R* is an uncompiled language, which can be used interactively. Using an uncompiled language helps empirical researchers, who may not have a deeper background in software engineering and its more abstract concepts to interactively test code. At the same time, using a packaged toolbox helps keep the code consistent. Chapter 6 aims to increase transparency for the researchers and contributes to bridging the gap between economics and information engineering by providing an interface between domain specific handling of time series objects with the *R* language and the popular general purpose relational DBMS *PostgreSQL*.

To sum up, despite long-known and new challenges and despite the advent of en vogue non-reactive data generating methods, surveys continue to be an attractive alternative to generate reliable, economic datasets. This is not only due the fact that survey researchers can build on a comprehensive strand of methodological literature but also due to new opportunities provided by technological advances, the inherent flexibility, and timeliness of surveys. Modern survey-based research can rely on a plethora of (online) tools that allow researchers to quickly reach out to potential participants and to flexibly adjust questionnaires according to recent political, social or economic events.

Chapter 2

No Answer Is an Answer - What Drives Unit Non-Response in Business Tendency Surveys?¹

¹This chapter is based on KOF Working Paper 363.

2.1 Introduction

In many countries researchers and policy makers alike rely on indicators built based on Business Tendency Survey (BTS) data.² Thus it is important for conductors of BTS to monitor the quality of their data closely and adapt their methodology to changes in participant behavior, as well as to new technical developments in information and communication technology. Within this context, survey researchers courting the attention of their increasingly distracted participants have to worry about non-response (Seiler, 2012). Researchers need to account for vast changes in participant behavior such as the degree of commitment towards requests in general. Strikingly, despite the idiosyncratic nature of BTS (Dillman, 2000), which hampers careless inference of results from other fields of survey based research – contributions of economic survey researchers to the methodological literature are rather scarce.

In this paper, we intend to contribute to the research on unit non-response by studying characteristics and attitudes of long-term BTS participants. We believe that the longitudinal structure and comparatively large samples of BTS contribute to methodological research beyond mere economic reasoning. While many studies on data quality use metrics at the survey-wave level, we are explicitly looking at the respondent-level. In other words, we are studying the relevance of participant characteristics to explain non-response as opposed to survey wave level effects. We pay particular attention to company characteristics such as company size, because, if these characteristics were unbalanced due to non-response, then economic inference from survey data would be influenced. We use a rich BTS dataset from the KOF Swiss Economic Institute and match it with data from an exclusive meta survey also conducted by the KOF. The matched dataset extends our set of firm characteristics with individual assessment and characteristics such as valuation of BTS, perceived response burden or participant gender.

The remainder of the paper is structured as follows: section 2.2 will give an overview of related studies in the field. In section 2.3 we describe our dataset and

²E.g. the *ifo Geschäftsklimaindex* (Germany), the *KOF Barometer* (Switzerland) or the Business Climate Indicator for the Euro area provided by the European Commission / DG ECFIN.

its different sources including a description of the meta survey. A detailed description of our empirical strategy follows in section 2.4. This methodology section consists of two major parts: First we explain how we account for selectivity issues caused by merging data from different surveys. Second, we discuss a hurdle model to model unit non-response count. Section 2.5 presents our estimation results. Finally section 2.6 summarizes our most robust results and provides an outlook for further research.

2.2 Literature

Despite the fact that Dillman (2000) points out the idiosyncrasies of surveying organizations: “people are asked to report information for an entity that is distinct from them personally”(Dillman, 2000, p. 324), and therefore justifies business survey-specific research, business survey specific contributions to survey methodology are comparatively scarce. As one of the few, Seiler (2013) studies non-response behavior based on German data from the ifo Institute’s manufacturing survey. Furthermore we have common ground with the work of Abberger et al. (2011) and Abberger et al. (2014) who gathered additional meta information by conducting surveys about surveys. Section 2.2.2 discusses meta information in BTS in greater detail. In addition to the BTS-specific literature this paper has been influenced by the work of Little and Rubin (2002) and Groves et al. (2002) on handling non-response and missing data in general. The ensuing subsection gives a brief introduction to non-response and missing data.

2.2.1 Non-response and missing data

The concept of non-response is usually defined as the rejection of an individual to participate in a survey (Groves et al., 2002). In order for non-response to occur some sort of contact between participant and survey conductor is assumed, e.g., in form of an email invitation. In this sense, individuals not included in the set of contacted individuals are not accounted for in the concept of non-response. However, in our paper the set of potential non-respondents is defined as a slightly larger group: Companies who were contacted but could not be reached are also counted

as non-respondents. Due to limitations in the dataset, we cannot distinguish these companies from those who were reached but explicitly refused to answer. However, this classification seems justifiable when we consider that the long-term customer relationship management of the KOF Swiss Economic Institute has filtered out letterbox companies and other irrelevant addresses over the years.

The literature distinguishes between item non-response and unit non-response. Item non-response describes a situation in which single questions are not answered whereas unit non-response means that a participant does not answer at all in a particular survey wave. In the latter case, the survey conductor does not gather any information on the company, which hinders sample adjustment within the post-processing of the dataset. In turn non-response potentially leads to several problems such as bias of the estimator's projection or inflation of its variance (Groves et al., 2002). This paper focuses on unit response and individual-level determinants of unit non-response in particular.

Several statistical methods deal with missing data caused by non-response. The ultimate success of these methods depends heavily on the systematics of non-response. Little and Rubin (2002) provide a useful classification of missing data: Data can be missing completely at random (MCAR), missing at random (MAR) or not be missing at random (NMAR). Data are called MCAR if the missing data mechanism depends on neither observed nor unobserved data. This is a very strict assumption in practice at times, which is why the weaker concept of MAR is widely used in applied research. MAR refers to a process in which the missing data mechanism depends only on observables. Data are said not to be missing at random (NMAR) if missingness depends on unobserved data.

In order to understand what causes non-response, a list of possible factors which have to be controlled for has to be found. Willimack et al. (2002) developed a conceptual framework for participation in business surveys.³ Taken from an earlier version of Groves and Couper (1998), Willimack et al. (2002) maintain the dichotomy

³The graphical depiction of the conceptual framework taken from Willimack et al. (2002) can be found in the appendix at the end of this thesis.

of factors, which are under the control of the researcher, and factors that the researcher has no influence on. The factors which are under control of the researcher are factors related to survey design, including sample design, instruments offered to answer the survey, time schedule and offered confidentiality. Factors, that are out of the researcher's control are divided into three subgroups, namely external environment, factors related to the business and factors related to the respondent. External factors include economic conditions and the legal and regulatory requirements. Among business related factors we find firm characteristics, such as size, industry, organizational structure as well as a company's philosophy and availability of resources. Finally we have a group containing factors of the respondent itself, such as sex, age, authority, capacity and motivation. Although the concept does not claim completeness, it offers a comprehensive list of factors which should be addressed when studying non-response.

2.2.2 Meta-information on BTS

The KOF BTS are conducted as multi-mode surveys where participants can take part either online or on paper. This opens up a plethora opportunities to gather meta information. Scholars of online research methods highlight the potential of non-reactive online research which analyzes subtly collected data (Reips, 2009) without the participant consciously answering questions. Such non-reactive research methods can also be combined with online surveys (Reips, 2009) resulting in so called paradata (data about the data generating process itself) and meta data (e.g. type of web browser or screen resolution) when data is collected through a web browser.

Though paradata promises insightful information about participant behavior that might eventually lead to non-response (Stieger and Reips, 2010), paradata in this strict sense cannot be used in our study. First, even for online participants paradata is not always available in desired quality for technical reasons, thus reliable paradata has not been collected for the entire sample of online participants. Using paradata would limit our research to sub-samples which would likely introduce selection problems. Yet we advocate the use of paradata and the introduction of tools to collect such information in BTS for future research. Second, a substantial

part of our participants continues to answer on paper. Hence, web paradata-based information would not be comparable across all participants.

Nevertheless, meta information on participants and the survey process itself is crucial to our work. Hence, we collected meta information using a multi-mode meta survey that allowed us to gather the same comparable information from participants of both modes: online and paper. Besides the methodological work of the scholars mentioned above, the pioneering work of Abberger et al. (2011), who conducted a survey about surveys, was inspirational to our work. While the study by Abberger et al. (2011) was limited to German trade firms we base our study on Abberger et al. (2014), which collected background information on firms from the Swiss service sector. The survey itself was conducted in the summer of 2013 and collected feedback data from all firms of BTS sample of Swiss service sector.

2.3 Data

All data used in this paper stem from surveys conducted by the KOF Swiss Economic Institute. KOF has implemented BTS in eight sectors, namely manufacturing, construction, retail trade, wholesale trade, service sector, hotels and restaurants, project engineering and financial and insurance activities. Today, KOF surveys include more than 11'000 companies. In addition to the regular BTS surveys, we conducted a feedback survey among the same set of participants as the KOF service sector survey. The following subsections introduce both datasets, the BTS data and the feedback data, used in this paper.

2.3.1 KOF BTS in the service sector

We chose to focus on the KOF BTS in the service sector because of its comparatively young age and decent share of online participants. Established only in 2006, paper based as well as online participation was possible from the first wave of the survey. Thus, participants were not potentially influenced by the introduction of an additional survey mode.

This paper uses a sample containing a total of 51'241 observations that were generated between the fourth quarter of 2006 and the fourth quarter of 2013. On average, close to 1'800 companies responded per wave during that timespan.

The available sample covers the following set of sectors according to the NACE 2-digit sector classification: 49-53, 58-63, 68-70,712, 72-75, 77-82, 86-88, 90-93, 95, 96. The sample does not contain firms of the accommodation and food service activities (NACE 55-56) sector, nor does it contain any firms of the financial and insurance sector (NACE 64-66). The survey in the service sector is conducted on a quarterly basis and the questionnaire consists of ten qualitative questions and one question with several non-mutually exclusive items.

2.3.2 A survey about surveys

In order to gather additional participant characteristics besides standard firm properties, we conducted a meta survey in summer 2013. We gain information on three major aspects:

- companies' understanding of the concept of the general business situation
- companies' valuation of the BTS
- measuring capacity utilization in the service sector

In advance of the survey we conducted interviewer pre-test to check the validity of our questions and items as well as the relevance of our set of questions to regular participants. We adjusted our questionnaire accordingly which left us with 10 questions covering the areas listed above.⁴

2.3.3 Merged sample

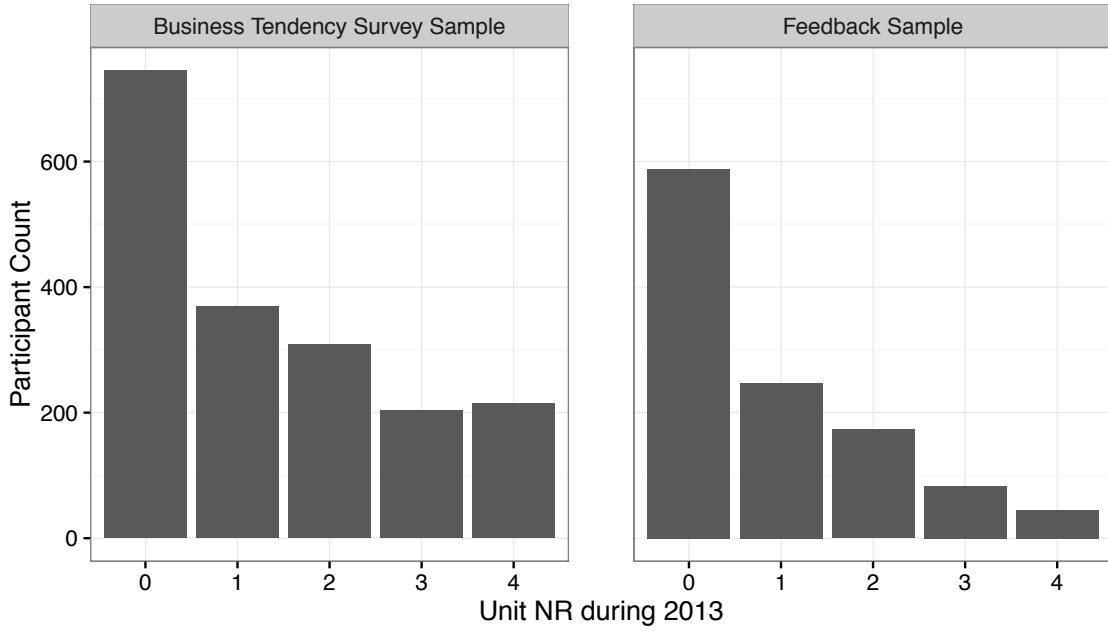
We merge the regular BTS sample and the sample of the feedback survey. Because the feedback survey was only conducted once we need to assume that the information

⁴The full questionnaire can be found in appendix to this chapter at the end of the dissertation.

gathered in this survey is time-invariant. This is a highly plausible assumption for a short period around the date of the feedback survey. The gathered information on company policies and structures is hardly volatile within a one year timespan. Hence, we decided to limit our non-response analysis to two quarters before and two quarters after the feedback survey. We divide our data into two subsets: data that was generated in the year of the feedback survey, namely 2013, and data that was gathered before.

We use the entire 2013 sample to analyze unit response in that year using standard firm characteristics and the additional firm data generated from the feedback survey. Any company that took part in the feedback survey conducted in the second quarter of 2013 and was invited to all waves of the regular BTS in 2013 is part of the sample. The 2013 subset contains 1136 companies. The data before 2013 is used to compute variables such as a company's survey track record, which is used to proxy a company's motivation to participate in 2013. Figure 2.1 shows the distribution of unit non-response count in 2013: in the left facet we see the distribution of unit non-response count for all companies that have been invited to take part in the regular BTS. The right facet shows the unit non-response count of the regular BTS in 2013 for those who also took part in the feedback survey. Section 2.4.1 describes how BTS data before 2013 is used to predict participation in the feedback sample in order to account for selectivity.

Figure 2.1: Unit Non-Response Count by Samples



2.3.4 Variables

This section presents all variables that are used in our analysis described in section 2.4. Table 2.1 classifies the displayed variables according to the conceptual framework (Willimack et al., 2002) discussed in section 2.2.

The first variable, *unit_nr_count*, is the number of quarters in which a firm did not respond over the course of 2013, i.e. the year of the feedback survey. The variable *unit_nr_count* is our variable of interest. All other variables will be used to explain the variation in *unit_nr_count*. The second row contains *unit_nr_ratio*, which captures the ratio of quarters in which a firm did not respond to the number of quarters since the firm's first invitation. The variable *unit_nr_ratio* can be regarded as a proxy for a firm's history of long-run motivation to take part in the survey before 2013. The variable in the third row, *avg_run*, is the average of a categorical timeliness indicator. A low average means that a company usually returns questionnaires immediately.

Table 2.1: Variable description

| Variable | Type | Dataset | Description |
|----------------------|--------------------|--|---|
| Company | | | |
| unit_nr_count | count | BTS | Number of unit non-response quarters during the year 2013 with zero if the company never missed a survey and four if the company did not participate at all. |
| unit_nr_ratio | continous | BTS | Record of past participation ratio: 0 = individual always participated before 2013, 1 = individual never participated before 2013. |
| avg_run | metric | BTS | Average of measured reaction time categories: 0 = fast, 1 = medium, 2 = slow. |
| valuation | metric | feedback | Importance of BTS: 1 = not important, 5 = important (5 items). |
| timeliness | categorical | feedback | Describes when a firm usually responds to a questionnaire. Answering straight away is the reference. time = whenever the person has time, deadline = answer right before the deadline, reminder = only when a reminder is received. |
| employees region | log categorical | company information company information | Size of a company: log number of employees. Espace, Lake Geneva Region, Northwestern Switzerland, Eastern Switzerland, Central Switzerland, Ticino and Zurich. |
| sector_class | categorical | company information | Part of the service sector, DL1 contains NACE 2 digit groups: 49-53, 58-63, DL2 contains 68-70, 72-75, 712, 77-82, and DL3 contains 86-88, 90-93, 95, 96. |
| capacity_utilization | continous | feedback | Capacity Utilization in percent, used to proxy business tendency at the firm level. |
| staff_shortage | binary | feedback | Indicate whether a company has problem hiring suitable staff. |
| Respondent | | | |
| gender | categorical | respondent information | Male, female, unknown (i.e no specific contact person). |
| language | categorical | respondent information | Italian, French, German. |
| position | categorical | feedback | Position of the person that answers the questionnaire: ownership or management, head of department, administration, someone else, no answer. |
| response_burden | binary | feedback | Perceived response burden high = 1, else 0. |
| surveys | count | feedback | Amount of surveys the respondents answer on average in one year. |
| Survey Design | | | |
| participation_mode | categorical | BTS | Paper, Online, Mixed. |

The variable *valuation* stems from the feedback survey and asks companies to evaluate the use of business tendency surveys in general. The variable *timeliness* captures a company's reported common reaction and thus proxies company's policy towards surveys. Companies were asked for their standards in processing regular BTS with respect to reaction time using the following items: answer right away, answer when time permits, answer before deadline or only answer after being reminded. It is difficult to form expectations on company policy, but we expect routines other than answering right away, to cause higher non-response. This is, because consistently rapid answers are more likely when a firm has incorporated the survey into their workflow routine. The variables *employees*, *region* and *sector_class* control for firm specific size, region and sector effects. In our view no clear expectations exist on the impact of different regions or sectors. The variables *capacity_utilization* and *staff_shortage* measure a company's availability of resources. Both variables are taken from the feedback survey. The variable *capacity_utilization* enters the model with indicated capacity utilization in percentage points and *staff_shortage* as a binary variable. We use both variables as an indicator for a firm's individual business situation. In the sense of Willimack et al. (2002), the variables presented above are company properties. In turn the variables described below relate to the respondent.

Ex-ante both variables, *gender* and *language*, are not expected to influence non-response in any direction. Besides male and female the gender variable has a third category labelled *unknown* which depicts cases in which companies did not provide a specific contact person. The work by Dillman (2000) shows that addressing a specific person within a company decreases non-response. The downside of addressing a specific person is the particular employee leaving the company. Several invitation letters were returned unopened with a remark that the addressed person has left the company. Both effects work in opposite directions and could cancel out in the aggregate. We expect the net effect to be negative, because fluctuation of employees within the sample is relatively small compared to the commitment effect of addressing a person directly. Therefore, addressing persons within a company directly is expected to decrease non-response. However, our dataset contains only 20 *unknown*

observations. Thus our investigation focuses on gender aspects as opposed to contact person vs. no-contact person set-ups.

The variable *position* indicates the position of the respondent within the firm. Both Abberger et al. (2011) and Abberger et al. (2014) show that business tendency surveys in Germany and Switzerland are mostly completed by persons in management positions. Tomaskovic-Devey et al. (1994) already pointed out that the position of the respondent can predict non-response. If the questionnaire is filled out by a person of a lower hierarchy level we expect non-response to increase for different reasons: Management level executives are likely to be more concerned with the company's reputation and are thus influenced by the commercial equivalent of social desirability. Furthermore, management level executives are more likely to be data savvy and might cherish additional information gained from business tendency surveys. Therefore we believe that management decisions to take part in surveys are more systematic than at lower hierarchy levels.

The fourth variable related to the respondent is a binary indicator on perceived response burden: If the perceived time to complete is longer than 10 minutes, it is labelled as burdensome. We expect the perception of burden to increase non-response. While the variable *response_burden* covers the burden caused by the KOF BTS in the service sector, the variable *surveys* counts the number of surveys from any survey conductor. Again both variables stem from the feedback survey.

The categorical variable *participation_mode* indicates the medium through which a company has chosen to participate. Participants can either take part on paper or online. Because we aggregate over one year the binary choice is extended by a mixed category indicating that not all quarters of 2013 were answered using the same participation mode. Ideally participation mode has no influence on non-response. Nevertheless, we expect an increase in non-response for online participants, as well as for those who switched the medium within 2013. Our expectations are based on the fact that companies who intend to drop out are sometimes convinced by KOF's customer relationship management that they can reduce the effort by answering online.

2.4 Methodology

Our dependent variable *unit_nr_count* can be suitably fitted using a Poisson distribution with excess zeros. We employ a two part hurdle model (Zeileis et al., 2008) to explain unit non-response count at the firm level. First, to measure firm-level unit non-response itself we count how many times a firm refused to answer the standard quarterly BTS during the course of 2013. Our estimations focus on the year 2013 because the feedback survey was only conducted once in summer 2013.

Unit non-response count $unit_nr_count \in \{0, 1, 2, 3, 4\}$ is the dependent variable in our model and we assume that the explanatory variables, which we obtained from the feedback survey, are time invariant over the course of one year. This assumption is highly plausible because variables such as the number of surveys answered per year or the responsible department are not likely to change over the course of one year unless company policy or the attitude of the management changes.⁵ The next section describes briefly how we account for selectivity issues. Subsection 2.4.2 describes how different count data models can be modelled in a GLM framework and subsection 2.4.3 continues to describe our main model.

2.4.1 Accounting for selectivity

As described in section 2.3.3 we make use of a merged sample, which contains data from the regular business tendency survey as well as from the feedback survey. Thus the sample used in our regressions contains only firms that are part of both datasets. This may lead to self-selection bias, as long-term, reliable participants of the regular surveys may be more likely to provide feedback than firms that frequently drop out from the regular surveys. In turn, we do not expect selection into our merged sample to occur completely at random. Thus we follow an approach similar to Little and Rubin (2002), who have shown how to account for selectivity when data is not missing completely at random. Figure 2.1 compares the distribution of our main variable of interest, namely unit non-response count across samples: especially the

⁵Please find the documentation of the questionnaire used in the feedback survey in the appendix at the end of the thesis.

tails indicate that the merged sample and the original BTS sample do not follow the same distribution.

We make use of inverse probability weighting to weigh (Horvitz and Thompson, 1952) our observations given their conditional propensity to end up in the merged sample as shown in equations 1 and 2. We use a standard probit model to estimate the probability \hat{p}_i of being in the merged sample – i.e. to be invited to the BTS and to have taken part in the feedback survey – given a set of firm characteristics.

$$w_i = \frac{1 - \hat{p}}{\hat{p}} \tag{2.1}$$

$$w_{norm,i} = \frac{w_i}{\frac{1}{n} \sum_{i=1}^n w_i} \tag{2.2}$$

These inverse probability weights $w_{norm,i}$ put more emphasis on observations that have a lower propensity to be in the merged sample given their characteristics. These observations would be under-represented without weighting. Remarkably, we were able to collect feedback data for several firms that did not take part in the regular survey all year. This information is very valuable since one usually does not have any information on mid- and long-term non-respondents.

2.4.2 Estimation of unit non-response in a GLM framework

Our analysis focuses on a hurdle model which is described in greater detail in section 2.4.3. The hurdle model is a two part model which consists of a Binomial and a Poisson component, both of which can be estimated in a GLM framework. Additionally, we estimate a standard Poisson model in section 2.5.2 to check the robustness of our results. All our regression models used to estimate non-response employ the same set of covariates X to explain the dependent variable $Y \in \{0, 1, 2, 3, 4\}$. Following Cox and Smith (1989), the hurdle model components and the models presented in the robustness section can be described as generalized linear models that consist of three components:

1. A *random component*

2. A *systematic component* η which is defined as: $\eta = \sum_1^k x_j \beta_i$ where k is the number of covariates
3. a link between the random and the systematic component which is defined as $\eta = g(\mu)$

Accordingly, $g(\cdot)$ is called the link function. Obviously, the discrete dependent variable $Y_i \in \{0, 1, 2, 3, 4\}$ is not represented well by the OLS model, i.e. the random component of Y was $N(\mu, \sigma^2)$ distributed and the *identity* function was used for the link. Thus, we configure the GLM framework presented above to describe a Poisson model, which provides a suitable benchmark for further computations and is also useful in our mixture model. Poisson models are among the most common models for count data and can also be specified within the GLM framework (Cameron and Trivedi, 2013). In the case of a Poisson model, the random component of Y is $P(\mu)$ distributed and the link function is $g(\cdot) = \log$. Cox and Smith (1989) and Cameron and Trivedi (2013) show that the Poisson model can also be estimated using iteratively weighted least squares (IWLS). Thus, our selectivity weights $w_{i,norm}$ from the previous section can be used in a straight forward manner. The Binomial part of the hurdle model is covered by the GLM framework in similar fashion as the Poisson model except that the random component is assumed to follow a Binomial distribution.

2.4.3 Estimation of unit non-response in a two part model

Figure 2.1 clearly shows that a unit non-response count of zero is the most prevalent among the sample firms in 2013. Though this is encouraging in the first place, an excess number of zeros does not comfortably suit a Poisson distribution. Hence, we estimate a two part model in order to account for these excess zeros. The binomial part models the decision whether to always answer the survey or drop out at least once. The Poisson part models the non-response behavior of firms which dropped out at least once. Further, Zeileis et al. (2008) suggest that the hurdle model might be slightly preferable over the related zero-inflation model, which is discussed in their same paper, because its interpretation is more intuitive. The distribution of our

dependent variable suggests that non-response is ultimately determined by multiple processes. In a first step, firms might decide whether to incorporate regular BTS into their workflow and thus answer by default or not. Firms that do not employ such a policy make the decision whether to take part in a survey wave upon receiving the questionnaire.

We specify a hurdle model which consists of two parts: First, a binomial part models the decision to participate in general as opposed to not participating. Second, a count part models how often non-response occurs $Y_i \in \{1, 2, 3, 4\}$ among those firms which dropped out at least once. Note that we use the same regressors $x_i = z_i$ to keep our models as comparable as possible though it is not required to use the same set of regressors. Formally:

$$f_{hurdle}(y; x, z, \beta, \gamma) = \begin{cases} f_{zero}(0; z, \gamma) & \text{if } y = 0 \\ (1 - f_{zero}(0; z, \gamma)) \cdot f_{count}(y; x, \beta) / (1 - f_{count}(0; x, \beta)) & \text{if } y > 0 \end{cases} \quad (2.3)$$

2.5 Discussion of estimation results

This section is divided into two subsections: First, we discuss the estimation results of our main two part hurdle model. Second, we present a robustness check which contains a weighted Poisson model as a benchmark, as well as models based on automatic variable selection procedures, such LASSO (Least Absolute Shrinkage and Selection Operator) and stepwise AIC.

2.5.1 Two part model estimation results

Table 2.2 shows the estimation results of our mixture model estimations. The left column lists the estimation results of the count data part and the right column shows the results for the binomial part of the hurdle model.

The non-response history of a company *unit_nr_ratio* is a significant indicator for future unit non-response in both parts of the model. The companies' reported

Table 2.2: Hurdle part

| | Count part | Binomial part |
|---------------------------|-----------------------------|-----------------------------|
| (Intercept) | -1.18 (0.43) ^{***} | -0.11 (0.79) |
| unit_nr_ratio | 0.97 (0.24) ^{***} | 2.65 (0.60) ^{***} |
| timeliness_time | 0.65 (0.20) ^{***} | 1.28 (0.22) ^{***} |
| timeliness_deadline | 0.58 (0.24) ^{**} | 0.73 (0.48) |
| timeliness_reminder | 0.99 (0.23) ^{***} | 2.74 (0.86) ^{***} |
| capacity_utilization | -0.00 (0.00) | -0.03 (0.01) ^{***} |
| position_department | 0.21 (0.10) ^{**} | -0.05 (0.32) |
| position_admin | -0.34 (0.25) | -0.01 (0.67) |
| position_other | 0.01 (0.22) | 1.10 (0.51) ^{**} |
| position_no_answer | 0.69 (0.19) ^{***} | -1.76 (0.83) ^{**} |
| response_burden_D | 0.29 (0.10) ^{***} | 0.46 (0.21) ^{**} |
| gender_female | 0.04 (0.11) | 0.47 (0.24) [*] |
| gender_unknown | 0.06 (0.26) | 0.94 (0.58) |
| participation_mode_online | 0.67 (0.19) ^{***} | 1.53 (0.25) ^{***} |
| participation_mode_mixed | 0.60 (0.30) ^{**} | 1.86 (0.91) ^{**} |
| avg_run | -0.08 (0.12) | 0.76 (0.36) ^{**} |
| valuation | -0.03 (0.03) | 0.05 (0.07) |
| surveys | -0.01 (0.01) | -0.04 (0.02) |
| staff_shortage | 0.05 (0.10) | -0.33 (0.22) |
| employees | -0.02 (0.04) | 0.06 (0.09) |
| region_Espace | 0.03 (0.17) | -0.45 (0.36) |
| region_LakeGeneva | -0.15 (0.29) | -0.64 (0.55) |
| region_Northwestern | -0.33 (0.23) | -0.25 (0.36) |
| region_Eastern | -0.48 (0.28) [*] | 0.36 (0.55) |
| region_Ticino | 1.08 (0.61) [*] | -4.46 (1.67) ^{***} |
| region_Central | -0.39 (0.17) ^{**} | 0.53 (0.54) |
| sector_class_DL2 | 0.18 (0.11) | 0.07 (0.28) |
| sector_class_DL3 | 0.07 (0.11) | -0.16 (0.25) |
| language_F | 0.38 (0.26) | 0.15 (0.53) |
| language_I | -0.97 (0.62) | 4.02 (1.52) ^{***} |
| AIC | 2122.67 | 2122.67 |
| Log Likelihood | -1001.34 | -1001.34 |
| Num. obs. | 865 | 865 |

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

common reaction time *timeliness* is a significant predictor of unit non-response. Every other behavior category increases unit non-response compared to answering

right away. Albeit it is noteworthy that companies which report to answer right before the deadline do not have a significantly lower propensity to not take part the entire year.

Though any significant influence of the business situation is clearly undesired, we find that firms which indicate higher *capacity_utilization* are more likely to employ a policy to generally return all waves of the KOF service sector survey. At first glance this finding seems to contradict Seiler (2013) and Harris-Kojetin and Tucker (1999), who find higher non-response during economically satisfactory times. However, their statement refers to inter-temporal comparison with the business cycle in mind, while we compare longterm participants with each other at the same point in the business cycle. We believe our results are plausible inasmuch that business success and responsible corporate citizenship, which includes participation in surveys, are both correlated with being well-organized and well-led. An upward bias, induced by the self-selection of well-managed companies into the sample, would be as undesired as any other bias. Identifying well led companies and adjusting the sample or post processing the results accordingly is a potential path to improve data quality. Clearly, more research, such as the regular analysis of different sectors and countries, is needed to monitor the influence of the business situation on non-response and to keep potential biases at bay.

When we account for excess zeros we see the effect of the position a respondent has within their company. We find that when surveys are answered outside of the management, administration and department heads and thus the ultimate decision whether to take part in general is made outside those parts of the company, companies are less likely to take part the entire year.⁶ The *position_other* category is not significant in the count part though. Whereas when the ultimate decision to take part is made by a department head instead of the management level, companies tend to not respond more often over the course of the year.

Perceived *response_burden* is a highly significant predictor of unit non-response. Accounting for excess zeros in fact emphasizes the effect of *response_burden* slightly.

⁶Only six companies chose to select *no_answer* in the position question of the feedback survey. Hence, we abstain from discussing possible reasons of the counterintuitive sign of this coefficient.

In particular, the effect on participation for a full year is stronger. Strikingly, we can see an effect of *gender*: females are significantly less likely to take part the entire year, though gender does not have a significant influence on the number of non-participation. This may be an indication that the *gender* of the participant, which is most likely at the management level, is actually correlated with a different kind of company culture, policy or age of the company. Yet, a gender effect is not expected and cannot be further explained by our data.

The results for *participation_mode* are in line with our expectations. Yet the effect of choosing not to answer the paper-based survey is emphasized in the binomial part and dampened in the count part. The true average reaction time *avg_run* is significant in the binomial part. Companies that return questionnaires later are on average less likely to have taken part in all waves of 2013. But there is no evidence that returning questionnaires later increases non-response count when we only consider firms that dropped out from the regular survey at least once during 2013.

Aside from capturing companies' general assessment of the value of BTS, *valuation* does not have a significant influence on non-response. This finding might be surprising at first but it is in line with findings for other types of surveys made in previous studies. Couper, M.P., Singer, E., Conrad, F.G., and Groves, R.M (2008) show that altruistic reasons are one of the most prominent motivations to take part in a survey. More than 30 percent of their participants indicated that they take part in surveys due to altruistic motives. Porst and von Biel (1995) report similar results for a study in the German-speaking area. Spitzmüller et al. (2007) also find that altruism is a driving force of response in organizational surveys. The number of *surveys* a company takes part in, *staff_shortage* and the size of the company (*employees*) do not play a significant role in either part of the hurdle model. Further, we cannot observe significant sectoral effects (*sector_class*).⁷

⁷Regional and language effects are difficult to interpret. Overall, we find that non-response count is not significantly influenced by region or language.

2.5.2 Variation and robustness of results

This section discusses other model variations in order to check the robustness of our results. While the discussion in the first subsection focused on specifications to find the distribution that suits our data best, this section compares our theoretical variable selection with two automatic variable selection procedures. Table 2.3 re-estimates our main model from table 2.2 using a weighted Poisson for benchmarking purposes. Column 2 compares the benchmark model to a stepwise AIC based variable selection (Venables and Ripley, 2002).⁸ Further, we check the robustness of our results using a shrinkage-based LASSO selection procedure, based on the *glmnet* implementation (Friedman et al., 2010). Note that we do not report shrunk coefficients and only use LASSO to suggest variables, which we include in a re-computed, weighted Poisson model.

Our main results, namely the effects of a firm’s response history *unit_nr_ratio*, the reported common reaction time *timeliness* and *participation_mode* remain significant for all selection methods. The LASSO-based selection drops the *response_burden* dummy. We can thus consider the result less robust, but should not overvalue this selection because it is not backed up by a theoretical basis. Also, *gender* is not selected based on the LASSO method which fuels previously stated thoughts that claiming a true gender effect is rather questionable and requires further research. Also, the undesired effect of *capacity_utilization* is left out in the selection mechanism presented in the third column. Although the fact that a variable, which is correlated with the business situation, is not selected comes in handy for the economic researcher, we should abstain from bold claims about the absence of a business situation-induced selection bias.

2.6 Conclusion and outlook

The process that determines non-response in BTS is fairly complex and cannot be modelled fully. Further complexity is introduced by the fact that participants are

⁸The *R* procedure we used runs both directions when running stepwise AIC-based selection.

Table 2.3: Automatically Selected Variables vs. Initial Model

| | Theoretical Selection | AIC | LASSO |
|---------------------------|-----------------------|-----------------|-----------------|
| (Intercept) | -0.99 (0.30)*** | -1.05 (0.30)*** | -1.70 (0.14)*** |
| avg_run | -0.04 (0.10) | | -0.07 (0.10) |
| unit_nr_ratio | 1.33 (0.19)*** | 1.23 (0.14)*** | 1.39 (0.21)*** |
| valuation | -0.01 (0.03) | | |
| timeliness_time | 0.85 (0.13)*** | 0.83 (0.13)*** | 0.89 (0.13)*** |
| timeliness_deadline | 0.63 (0.20)*** | 0.64 (0.19)*** | 0.74 (0.19)*** |
| timeliness_reminder | 1.21 (0.17)*** | 1.24 (0.19)*** | 1.33 (0.18)*** |
| employees | 0.02 (0.03) | | |
| region_Espace | -0.08 (0.16) | | |
| region_LakeGeneva | -0.35 (0.28) | | |
| region_Northwestern | -0.23 (0.15) | | |
| region_Eastern | -0.24 (0.14)* | | |
| region_Ticino | -1.07 (0.28)*** | | |
| region_Central | -0.05 (0.12) | | |
| sector_class_DL2 | 0.15 (0.10) | | |
| sector_class_DL3 | 0.06 (0.09) | | |
| capacity_utilization | -0.01 (0.00)*** | -0.01 (0.00)** | |
| staff_shortage | -0.09 (0.09) | | |
| language_F | 0.37 (0.26) | 0.15 (0.12) | 0.18 (0.12) |
| language_I | 1.10 (0.23)*** | 0.18 (0.16) | 0.11 (0.16) |
| position_department | 0.16 (0.10) | 0.17 (0.10)* | |
| position_admin | -0.20 (0.19) | -0.17 (0.18) | |
| position_other | 0.19 (0.16) | 0.22 (0.14) | |
| position_no_answer | -0.17 (0.61) | -0.13 (0.69) | |
| response_burden_D | 0.25 (0.08)*** | 0.19 (0.08)** | |
| surveys | -0.01 (0.01) | -0.01 (0.01) | |
| gender_female | 0.15 (0.09)* | 0.23 (0.08)*** | |
| gender_unknown | 0.32 (0.19)* | 0.27 (0.18) | |
| participation_mode_online | 0.85 (0.13)*** | 0.81 (0.13)*** | 0.78 (0.13)*** |
| participation_mode_mixed | 0.74 (0.25)*** | 0.81 (0.27)*** | 0.71 (0.25)*** |
| AIC | 2139.67 | 2131.26 | 2150.31 |
| BIC | 2282.56 | 2216.99 | 2197.94 |
| Log Likelihood | -1039.84 | -1047.63 | -1065.15 |
| Deviance | 775.50 | 791.09 | 826.13 |
| Num. obs. | 865 | 865 | 865 |

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

answering on behalf of companies instead of representing themselves. Nevertheless, our regressions are able to model a substantial part of the process described in the conceptual framework of Willimack et al. (2002). In our paper, we are able to

identify important drivers of unit non-response and show meta information, which has so far largely been neglected, explains a substantial amount of unit non-response. We found that a company's response history, the reported common reaction time and the participation mode to be important predictors of unit non-response in all of our model set-ups. We believe that our mode effect is not necessarily a true mode effect, as presenting questions online does not affect the result or the likelihood to respond. Companies which are less committed and answer more ad hoc are rather induced to answer given the ease of an online survey. Also, companies which drop out of the paper-based survey are encouraged to reduce the burden by participating online. Thus, our results are not based on a pure medium effect. In addition, we consider perceived response burden to be among the strongest indicators of unit non-response. Hence, we advise conductors of BTS to put emphasis on the optimization of usability in order to reduce the perceived response burden.

By measuring non-response as the cumulative non-response count of a company in 2013, we are able to model this unit non-response count via a Poisson distribution with excess zeros. This suggests that unit response in regular BTS cannot be explained with one single data generating process, but is rather a mixture of processes. The first decision that a company makes with respect to surveys might be whether to include regular survey participation into their usual workflow. If a company does not have a policy of participation, the decision to participate is made during every wave of the survey. It is plausible that these two different decisions are driven by different factors.

Fortunately, we do not find an effect of company size and sector affiliation on the propensity to not respond. This is a desirable result from an economic point of view, as there is no evidence of self-selection of larger firms or particular sectors. Further, we find no evidence of an effect of valuation of BTS or the number of surveys a company answers on the propensity not to respond. This is only mildly surprising as it is in line with previous findings in the literature, which state that altruism is a main driver for survey participation (Couper, M.P., Singer, E., Conrad, F.G., and Groves, R.M, 2008).

However, we find that high capacity utilization – reported through the one time feedback survey – reduces unit response. On the one hand this finding may indicate that our sector controls are not sufficient to account for firm-level heterogeneity, on the other hand, this finding may well be driven by the undeniable correlation between capacity utilization and the business situation. This latter implication is concerning from an economic standpoint: companies in a better business situation – proxied by capacity utilization – tend to answer more regularly.

Female participants seem to be less likely to respond regularly. Though this an intriguing finding, we abstain from claiming a true gender effect. We find gender to be significant only in the binomial part of our two part model, suggesting that female participants were less likely to answer all waves in 2013, which might be driven by company characteristics that are correlated with the prevalence of women in executive positions. Company characteristics that are correlated with not including surveys in the regular workflow and employing female executives may rather drive this finding.

We showed that meta information such as the response burden or gender improve modelling unit non-response in BTS. Furthermore, our results show that two groups of firms with different non-response behavior exist. Further feedback data from other sectors is needed to generalize our findings. We suggest implementing regular feedback surveys that alternate through sectors in different periods to gather meta information from each sector every couple of years. This would help monitoring changes in the data generating process more closely and help improving methodology in economic surveys and to adapting surveys and panel management. Our work also emphasizes the importance of participant relationship management. If participants do not drop out at random they need to be replaced carefully. Also our results with respect to company policies and overall participation patterns emphasize the importance of recruiting new companies explicitly to a panel survey, as opposed to recruiting firms to a single wave.

Chapter 3

Macro and Micro Level Impulse Responses: A Survey Experimental Identification Procedure¹

¹This chapter is based on KOF Working Paper 386.

3.1 Introduction

While aggregate economic outcomes are a composite of individual business activities, firms' decisions in response to macroeconomic shocks are highly complex and manifold. Often the resulting responses at the aggregate level turn out to be ambiguous, hampering the analysis of causal effects. In these situations, identifying the effects of macroeconomic shocks already at the firm level could be fruitful. However, this is a formidable task given the prevailing non-experimental data available to macroeconomists.

In this paper, we investigate the determinants of responses to macroeconomic shocks on the macro and micro level using a survey experimental procedure.² We show that the effects identified with our survey experiment are conceptually equivalent to impulse responses obtained from vector autoregressions (VARs). However, in contrast to VARs, there is no need to impose parametrical restrictions to identify exogenous macroeconomic shocks. The survey-based approach does also not require extensive longitudinal data, which can be subject to structural breaks in the data generating process but sets up an on the spot analysis by design. Instead of indirectly deducting causal effects from historical time series, we retrieve such responses directly by asking firms.

We employ the survey experimental approach to study the effects of oil price shocks at different levels of disaggregation. Oil price shocks are among the most prominent and important macroeconomic shocks (see, e.g., Hamilton, 1988 and Kilian, 2009). Responses to oil price shocks can lead to inflationary pressure and may have real effects on the economy. However, little is known about the effects of oil price shocks at the firm level. We contribute to the existing literature by conducting a survey experiment to identify macroeconomic shocks without any parametric restrictions and thereby generating a revision-free firm-level dataset. More precisely,

²The expression of survey experiments originates from the field of psychology (see reviews by Sniderman and Grob, 1996, Gilens, 2002, and Guterbock and Nock, 2010). See Kuziemko et al. (2015) for a survey experiment on income inequality and Drechsel et al. (2015a) for a survey experiment on exchange rate shocks.

we attached a special questionnaire to the July 2012 wave of the KOF Swiss Economic Institute Investment Survey, a major statistical survey on firms' financial outcomes and plans in Switzerland. We mimic the official macroeconomic data generation procedure of statistical agencies (as would be the case for GDP): In a first step, we collect firm-level data from a representative sample of firms in the Swiss economy. 1037 Swiss firms completed our questionnaire, which asked for changes in firms' turnover and sales prices in response to an exogenous shock to the oil price. Thereafter, we aggregate the micro-level data to the industry- and economy-wide level using standard national accounting procedures for constructing of macro time series from micro data.³

We compare impulse responses derived from survey experiments to results obtained from a structural Bayesian VAR (BVAR).⁴ We find that the BVAR analysis involving standard identifying assumptions conforms to the results obtained from the experimental survey set-up that works without imposing any identifying restrictions. We observe a lot of heterogeneity in impulse responses. In an industry-level exercise similar to the approach undertaken in Lee and Ni (2002), we distinguish whether industries' impulse responses to an oil supply shock are dominated by cost-push or the demand channels. We find that in more oil consuming industries cost-push channel is dominant. In contrast, for less oil consuming industries the demand channel tends to prevail.

We also analyze possible driving forces of the causal effects by regressing firm-level impulse responses on a set of covariates such as oil consumption, market power, firm size, industry membership and time effects. We find that market power and firm size do not explain a significant part of firm level responses to an oil price shock, whereas oil intensity is an important and statistically significant influencing factor. We use our firm-level data to disentangle industry membership effects from oil consumption. At the industry level, we find little correlation between oil consumption

³See, e.g., European Commission, 2007.

⁴To achieve such comparability, a number of conditions have to be met. The framework of the survey experiment should specify the magnitude of the macroeconomic shock, measure the response on a ratio scale, and depict the timing of the effect. Further, a survey experiment should be conducted among a representative sample of firms.

and output responsiveness to an oil price shock, confirming the results in Lee and Ni (2002). However, when controlling for industry membership at the firm level, we do find a significant effect of firms' oil consumption on output reactions, thereby solving an existing puzzle in the literature.

Furthermore, we show that the usual identification assumptions in VARs imply restrictions on impulse responses. Depending on the prevalent industry structure and magnitude of the underlying firm-level response, such imposed restrictions might bias. We suggest an easy way to test for such a bias using *structural* micro data obtained from our survey experiment.

The remainder of the paper is structured as follows. The following section 3.2 discusses impulse responses based on survey experiments, shows the linkages between firm level and aggregate impulse responses, and elaborates on survey validity. Section 3.3 describes the survey setup and experimental design to identify the effects of oil price shocks. Section 4.3 presents the results obtained from our oil shock macro survey experiment and section 4.4 concludes. Technical derivations and more detailed results appear in the appendix.

3.2 Method

This section describes how survey experiments can be used to measure causal effects of macroeconomic shocks on a firm-level basis. Aggregate causal effects are commonly of interest when studying the responses to macroeconomic shocks. Yet, aggregate data are a composite of firm-level observations. If one truly aims to understand the driving forces behind economic reactions to macroeconomic shocks, it is helpful to study responses at the firm level.

In the following subsection, we explain our procedure to generate firm-level impulse responses to macroeconomic shocks using survey experiments and highlight several virtues of the approach. Subsequently, we discuss the mapping from firm-level impulse responses to aggregate impulse responses. This analysis sheds light on the implications of the aggregate restrictions commonly applied in VARs for the distribution of firm-level impulse responses. And lastly, we focus on the process of

forming expectations under a scenario and discuss the validity of our survey experiment.

3.2.1 Generating impulse responses from survey experiments

Survey-based impulse responses are generated by mimicking official data collection procedures among representative samples of firms. The following three features help to achieve representative and valid outcomes: i) The survey comes from a respected institution, or the statistical agency itself or an affiliated agency. ii) It is conducted on a regular basis (though the actual questionnaire content may change over time). And iii), respondents come from the top management of firms.

In a first step, we collect firms' expectations (or projections) on key financial figures such as their revenues, expenses and investments. These expectations are usually recorded in firms' business information systems and can be easily retrieved. Next, we confront firm respondents with a macroeconomic shock scenario, all other things being equal, and ask them to report their expectations under this scenario.⁵ The difference between the expectations in the baseline scenario (= control scenario) and in the shock scenario (= treatment scenario) gives the expected firm-specific causal effect (= treatment effect) of the shock. In a last step, we aggregate the expected firm-specific causal effects over firms to get the expected macroeconomic causal effect of the shock. The aggregation is done with standard procedures used in statistical agencies to build macroeconomic series from micro data.

More formally, the set-up can be expressed for firm $i = 1, \dots, I$, period $t = 1, \dots, T$ and horizon $s = 1, \dots, S$ as

$$\psi_{i,t+s} = E_{i,t}(y_{i,t+s}|\eta_{i,t} = 1) - E_{i,t}(y_{i,t+s}|\eta_{i,t} = 0), \quad (3.1)$$

and

$$\psi_{t+s} = \sum_{i=1}^I \omega_i \psi_{i,t+s}, \quad (3.2)$$

⁵It is important to phrase the scenario in a way that firm respondents can understand. One cannot assume that respondents understand the word shock in the same way as economists.

where $\psi_{i,t+s}$ is the dynamic expected causal effect for firm i , $E_{i,t}(y_{i,t+s})$ is the firm-specific expectation for horizon s at period t , $\eta_{i,t}$ is the treatment variable with $\eta_{i,t} = 1$ when firm i receives the shock treatment and $\eta_{i,t} = 0$ otherwise, ψ_{t+s} is the dynamic expected macroeconomic causal effect of the shock, and ω_i is the aggregation weight of firm i .⁶

The dynamic causal effect described in (3.1) and (3.2) is equivalent to the definition of impulse responses in the time series literature, where the treatment is a shock at time t with its effects s periods after the shock has occurred (e.g., Hamilton, 1994). Given this definitional equivalence, we refer to the above equations as the firm-level or macroeconomic *survey-based impulse response* to the shock.

Survey-based impulse responses have several virtues. First, survey-based impulse responses provide a convenient way to identify macroeconomic shocks through its experimental design. Indeed, these shocks do not need to be identified econometrically. They can be extracted by confronting firm executives with a scenario in which their firm is hit by a macroeconomic shock. Thus, survey-based impulse responses can be used to test macroeconomic theories without presuming any prior economic theory.⁷

A second strength of survey-based impulse responses is their bottom-up structure. They allow the study of heterogeneous causal effects at the firm level with relatively modest requirements.⁸ By employing an appropriate firm weighting scheme, the firm-level impulse responses can be aggregated to the industry, sector or economy-wide level. Thus, impulse responses based on surveys allow for analysis at any desired level of (dis-)aggregation.

⁶By asking the *same* firm representative about a control scenario and a treatment scenario, the set-up follows a within-subject design (see Charness et al., 2012). One might also follow a between-subjects design by randomly assigning different scenarios to firms. Both within-subject and between-subjects designs allow for *multiple* treatment scenarios (effect of shock A, effect of shock B, effect of shock C, ...).

⁷See the discussion in Angrist and Pischke (2010) and the replies by, e.g., Leamer (2010), Sims (2010), and Stock (2010).

⁸Data availability often poses serious obstacles for empirical work on a firm-level basis. Survey experiments overcome this issue as they create tailored micro datasets for the question of interest.

Third, survey-based impulse responses measure the effect of a shock at the time the survey was conducted. As a consequence, policy makers can make use of this procedure to determine the effects of shocks “on the spot”, for instance, at times of a suspected structural break.⁹

Survey-based impulse responses have yet another virtue: one can test for possible biases in aggregate, identifying restrictions. The argument is further exposed in the next subsection.

3.2.2 Firm level vs. aggregate impulse responses

The aggregate level and the firm level are linked via the employed aggregation procedure. If the aggregation process is known, aggregate impulse responses to shocks can be decomposed into firm-level impulse responses. A popular workhorse model for studying causal effects on the aggregate level are VARs. In order to determine causal effects from VAR analysis, identifying assumptions have to be set. Given the connection between aggregate and firm-level data, restrictions at the aggregate level imply certain restrictions on the distribution of firm-level impulse responses.

If detailed time series information were available at the firm level, a VAR could be set up for each firm, which obtains the dynamic causal effects of shocks at different aggregation levels on firms’ key variables, say, output and prices. The firm level VAR for firm i for $i = 1, \dots, I$ at time t can therefore be expressed as follows:

$$y_{i,t} = c_i + \Phi_{i,1}y_{i,t-1} + \Phi_{i,2}y_{i,t-2} + \dots + \Phi_{i,p}y_{i,t-p} + u_{i,t}, \quad (3.3)$$

where $y_{i,t}$ is $m \times 1$, c_i is a $m \times 1$ vector of intercepts, $\Phi_{i,p}$ is a $m \times m$ coefficient matrix and $u_{i,t}$ is an $m \times 1$ vector of disturbances with zero mean and variance covariance matrix Σ_i . The reduced form error $u_{i,t}$ can also be written as a linear combination of structural shocks $\epsilon_{i,t}$, hence $u_{i,t} = A_i^{-1}\epsilon_{i,t}$, where A_i is a $m \times m$

⁹Further, by conditioning on the actual information set of economic agents, the scope of the information set is not an issue for survey-based impulse responses (see Rudebusch, 1998). In addition, survey-based impulse responses involve no assumptions about the expectation formation of economic agents.

nonsingular coefficient matrix and $\epsilon_{i,t} \sim N(0, I)$. The VAR(p) model in its VAR(1) companion form then is:

$$Y_{i,t} = C_i + F_i Y_{i,t-1} + U_{i,t}, \quad (3.4)$$

where $Y_{i,t} = [y'_{i,t}, y'_{i,t-1}, \dots, y'_{i,t-p+1}]'$, $U_{i,t} = [u'_{i,t}, 0, \dots, 0]'$, and F_i is the $mp \times mp$ companion matrix containing the VAR coefficients with

$$F_i = \begin{bmatrix} \Phi_{i,1} & \Phi_{i,2} & \cdots & \Phi_{i,p-1} & \Phi_{i,p} \\ I_m & 0 & \cdots & 0 & 0 \\ 0 & I_m & 0 & 0 & 0 \\ \vdots & & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & I_m & 0 \end{bmatrix}$$

where I_m is an $m \times m$ identity matrix. Given that the VAR follows a stable process, (3.3) can be rewritten in its infinitely moving average representation

$$Y_{i,t} = \mu_i + \sum_{s=0}^{\infty} F_i^s U_{i,t-s}, \quad (3.5)$$

where $\mu_i = (I_{mp} - F_i)^{-1} C_i$.¹⁰ (3.5) can then also be written in its more condensed form

$$y_{i,t} = J\mu_i + \sum_{s=0}^{\infty} \Psi_{i,s} \epsilon_{i,t-s}. \quad (3.6)$$

where $J = [I_m, 0, \dots, 0]$ is $m \times mp$ and $\Psi_{i,s} = JF_i^s J' A_i^{-1}$ is the firm-specific impulse response matrix at horizon s to a shock in $\epsilon_{i,t}$, i.e. $\Psi_{i,s} = \frac{\partial y_{i,t+s}}{\partial \epsilon_{i,t}}$. Let us assume that we are only interested in aggregate shocks, i.e. in ϵ_t .¹¹ The impulse response function in this case implies $\Psi_{i,s} = \frac{\partial y_{i,t+s}}{\partial \epsilon_t}$, which is equal to the definition of survey based impulse responses discussed in section 3.2.1.

¹⁰See, e.g., Hamilton (1994) or Lütkepohl (2005).

¹¹The aggregate shock can also be interpreted as a common shock to all firms. This could be implemented using factor analysis, decomposing the structural errors $\epsilon_{i,t}$ into a common, an idiosyncratic and a firm-specific part.

Consider now that we aim to aggregate these firm-specific variables in $y_{i,t}$ using standard procedures.¹² The aggregate time series created will be a convex combination of the underlying disaggregated series:

$$y_t = \sum_{i=1}^I \omega_i y_{i,t}, \quad (3.7)$$

where $\omega_i \geq 0$ and $\sum_{i=1}^I \omega_i = 1$. Substituting (3.6) into (3.7) results in the following expression

$$y_t = \sum_{i=1}^I \omega_i (J\mu_i + \sum_{s=0}^{\infty} \Psi_{i,s} \epsilon_{t-s}), \quad (3.8)$$

with the implied aggregated impulse response matrix $\Psi_s = \sum_{i=1}^I \omega_i \Psi_{i,s}$ at horizon s to a shock in ϵ_t , i.e. $\Psi_s = \frac{\partial y_{t+s}}{\partial \epsilon_t}$.

Similar to the firm-level VAR described above, the aggregate VAR with aggregated variables can be expressed as follows

$$y_t = c + \Phi_1 y_{t-1} + \Phi_2 y_{t-2} + \dots + \Phi_p y_{t-p} + u_t \quad (3.9)$$

with its infinitely moving average representation as

$$y_t = J\mu + \sum_{s=0}^{\infty} \Psi_s \epsilon_{t-s}, \quad (3.10)$$

where the general structure is similar to the one described for the firm-level VAR, except for the parameters, which are now no longer firm specific. The aggregate impulse response matrix at horizon s to a shock in ϵ_t hence is $\Psi_s = JF^s J'A^{-1}$, implying again $\Psi_s = \frac{\partial y_{t+s}}{\partial \epsilon_t}$. From (3.8) it follows that

$$\Psi_s = \frac{\partial y_{t+s}}{\partial \epsilon_t} = \sum_{i=1}^I \omega_i \Psi_{i,s}. \quad (3.11)$$

¹²See, e.g., European Commission, 2007.

Identifying structural shocks in the aggregated VAR described in (3.9) requires certain restrictions on the parameter space. Usually, either exact or sign restrictions are imposed on the impulse response matrix Ψ_s .¹³ In contrast to VARs, generating survey-based impulse responses and hence obtaining firm-specific impulse responses does not require parametric restrictions.

Define $\psi_{kj,s}$ as the response of the k -th variable in y_t to the j -th shock at horizon s . Assume now that the identification of the system involves a parametric restriction and that this restriction is imposed on the response of the k -th variable to the j -th shock at horizon h . Without loss of generality, we assume that the restriction can be expressed as $\psi_{kj,h} = r$. Using the relationship between aggregated impulse responses and disaggregated impulse responses in (3.11), we obtain the following expression for the restriction imposed:

$$\sum_{i=1}^I \omega_i \psi_{kj,h}^i = r, \quad (3.12)$$

where $\psi_{kj,h}^i$ is the response of the k -th variable to the j -th at the horizon h of firm i . Hence, any deviation of r from the aggregated impulse response function $\sum_{i=1}^I \omega_i \psi_{kj,h}^i$ results in a bias. Consider the following quadratic loss function for the restricted elements in Ψ_h

$$\left(r - \sum_{i=1}^I \omega_i \psi_{kj,h}^i\right)^2. \quad (3.13)$$

The loss function implies that the loss is zero when the different firm-specific impulse responses can be combined in way such that they sum exactly to r . More importantly, the degree of the bias imposed by the aggregate restriction depends crucially on the weight of these biased firms' impulse response functions and the magnitude of these impulse responses.

Firm-specific impulse responses and survey-based impulse responses, respectively, allow testing for the existence of a bias in aggregate restrictions, occurring

¹³Zero restrictions on Ψ_0 are similar to restrictions used in Sims and Zha (1998) to identify an SVAR. In case Ψ_0 is a triangular matrix the Cholesky decomposition of Σ is of the form $\Sigma = \Psi_0 \Psi_0'$. It is also possible to impose sign restrictions on Ψ_s for different s as described in, e.g., Faust (1998) and Uhlig (2005).

for instance in VARs. This test could include exclusion restrictions, where r corresponds to a specific value or sign restrictions, where r corresponds to the positive and negative region, respectively. Such a test resembles a two-sided student t-test for exclusion restrictions (e.g. Cholesky decomposition, $r = 0$) and a one-sided t-test in case of sign restrictions, where $\sum_{i=1}^I \omega_i \hat{\psi}_{kj,h}^i = r$ is the relationship to be tested with $\hat{\psi}$ denoting the impulse responses obtained from our survey experiment. The results are reported in Table 3.3).

3.2.3 Validity of the survey experiment

This section takes a closer look at the process of forming expectations for both the baseline and the scenario. In the case of our survey, firms form expectations on operating figures under the current situation as well as under the assumption of an exogenous shock. We can write the model process of forming baseline expectations as:

$$E_{i,t}(y_{i,t+s}) = \sum_{k=0}^K g(y_{i,t-k}) + \sum_{j=1}^J p_{i,j} f_{i,j}(\pi_{j,t}), \quad (3.14)$$

with $\sum_{j=1}^J p_j = 1$, where $g()$ is a function of firm i -th lagged K operating figures and $f()$ is a function of J possible states of a firm's environment.

π_j denotes the j -th state a firm expects to materialize in the future. Then p_j is the probability that the j -th state actually materializes. The states π capture all of the firms belief about the future and are thus mutually exclusive.

Given that companies are able to form expectations, which are consistent with realizations of their operating figures, we are confident that firms are able to assess a ceteris paribus scenario which is much less complex than the baseline situation described above. When setting an exogenous shock and explicitly ruling out other events, a firm is locked in on a single scenario π_z such that the expectation formulation is exogenously changed with $p_{j=z} = 1$ and $p_{j \neq z} = 0$. Hence, the process of forming expectations under the scenario of exogenous shock can be simplified as:

$$E_{i,t}(y_{i,t+s}) = \sum_{k=0}^K g(y_{i,t-k}) + f_{i,z}(\pi_z) \quad (3.15)$$

This process is arguably much simpler as participants do not have to build expectations for multiple states, nor do they have to assess the probability of their occurrence. Further, only a single functional form $f_{i,z}(\pi_z)$ needs to be considered. Hence, we conclude that if firms are able to form baseline expectations that are in line with their baseline realizations, firms should be able to form reasonable expectations under a much simpler scenario. In order for this conclusion to be made, it is important that the the treatment scenario must be realistic in the sense that respondents have been confronted with similar scenarios in the past and/or that they have previously considered the scenario and its effects (see Gaines et al. (2006)).¹⁴ Thus, we believe that if firms' forecasting errors with respect to the realizations of their baseline scenarios are on average zero, firms should also be able to make conditional forecasts.¹⁵

3.3 An application to oil price shocks

While the previous section described the procedure to generate survey-based impulse responses, this section presents an application in the form of oil price shocks.¹⁶ Oil price shocks are well-suited to illustrate survey-based impulse responses. They are easily understood and imagined for survey respondents and have frequently occurred in the past.


The development of the oil price can be seen in figure 3.1. Since the late 1990s, the oil price in Swiss Francs (CHF) has been fairly volatile around an upward trend.

¹⁴We consider this condition to be fulfilled in our application in Section 3.3. This notwithstanding, economists who believe that economic agents build rational expectations should generally have no reason to mistrust impulse responses generated from survey experiments.

¹⁵The appendix to this chapter at the end of the thesis provides a robustness check on this particular assumption. Further robustness checks, such as non-response analysis, are also in the appendix at the end of the thesis.

¹⁶see Hamilton (2008) and Kilian (2008) for comprehensive reviews of the literature.

Figure 3.1: Price of crude oil in CHF per barrel



Prior to the recent financial crisis, oil prices have peaked at more than 140 CHF per barrel Brent. In the aftermath of the financial crisis, the oil price dropped by more than 50%, with a swift recovery afterwards. During the winter 2014/15 and spring 2015, prices collapsed again.

3.3.1 Survey set-up

Our data stem from a questionnaire attached to the semi-annual KOF Swiss Economic Institute Investment Survey during the summer 2012 wave.¹⁷ The characteristics of the underlying sample are representative of the Swiss economy. Detailed information on the sampling procedure can be found in appendix to this chapter at the end of the thesis.

¹⁷The survey was conducted as a multi-mode survey: Part of the survey has been paper-based, the other part was conducted online using a self-hosted instance of LimeSurvey (www.limesurvey.org).

1037 Swiss firms completed the additional set of questions of which 85 firms are from the construction, 434 from manufacturing and 518 from services sector. Firms' responses to KOF surveys come mostly from CEOs and CFOs.¹⁸ Respondents are taking part in KOF enterprise surveys on a regular basis and are accustomed to KOF questionnaire design. In order to ensure relevance of our questions to practitioners, we conducted an interviewer pre-test among a group of selected firms and adjusted our questionnaire according to the feedback.

Participating firms received an invitation letter and the questionnaires in paper and electronic format in order to facilitate participation. Anonymity of responses has been guaranteed.¹⁹ If participants did not respond within 18 days they received a reminder. Firms that did not participate after the initial reminder, received an additional telephone reminder after an additional two weeks. Questionnaires were sent out in German, French and Italian according to firms' preferences.

3.3.2 Experimental design

This section explains the design of our macro survey experiment. Prior to posing scenario questions, participants stated their key financial figures for the past (2010, 2011), present (1st half of 2012) and predictions for the future (2nd half of 2012, 2013). This step is helpful in setting the benchmark for the survey experiment.

The questionnaire then confronts participants with the counterfactual situation of an oil price shock and asks them to re-evaluate their answers under the following hypothetical scenario:

Suppose, the oil price increases by 30% within the next month under otherwise constant economic circumstances and will remain 30% above your previous expectations of the oil price development. Please indicate how your financial figures would change compared to your previous expectations for these figures.²⁰

¹⁸See, e.g., Abberger et al. (2014).

¹⁹All KOF surveys are subject to Swiss statistics law.

²⁰The base level for the oil price shock is the price of oil shortly before sending out the survey. The jump in the oil price is set to one standard deviation of the monthly price series (barrel crude

The answer options for total turnover are as follows (the complete questionnaire can be found in the appendix of this chapter at the end of this thesis.):

| <u>2nd Term 2012</u> | | | | | | | | | | | |
|---------------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|-------------------------------------|
| ≤-7.5% | -5% | -3% | -2% | -1% | 0% | 1% | 2% | 3% | 5% | ≥7.5% | N/A |
| <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> |

| <u>2013</u> | | | | | | | | | | | |
|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|-------------------------------------|
| ≤-7.5% | -5% | -3% | -2% | -1% | 0% | 1% | 2% | 3% | 5% | ≥7.5% | N/A |
| <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> |

It is important to note that the above scenario question is designed such that answers equal the firm-specific dynamic causal effect of the oil price shock (“treatment effect”). Accordingly, (3.1) can be specified as

$$\psi_{i,t+s} = E_{i,t}(y_{i,t+s}|\eta_t = 1) - E_{i,t}(y_{i,t+s}|\eta_t = 0), \quad \text{for } s = 6 \text{ months, 18 months}$$

where $E_{i,t}(y_{i,t+s}|\eta_t = 1)$ is firm i 's expected turnover at horizon s given that the oil price shock occurred at time t and $E_{i,t}(y_{i,t+s}|\eta_t = 0)$ is its expected turnover at horizon s given that the oil price shock did not occur, ceteris paribus.

In the same style, the questionnaire asked participants to evaluate the effect of the oil price shock on average purchase prices, total expenditure, average domestic sale prices and average foreign sale prices for the second half of 2012 (after 6 months) and for 2013 (after 18 months). The questionnaire further asked for firms' (pre-shock) expenditure on oil products in terms of total expenses (“oil share”), exports in terms of total turnover (“export share”) and imports in terms of total expenses (“import share”). The appendix at the end of the thesis gives a comprehensive list of variables used in this study. The defined shock is similar to a permanent cost-push shock. We exclude (global or Swiss wide) oil demand shocks by stating that firms should consider their responses “under otherwise constant economic circumstances”. The shock is specified as a permanent shock by stating that “the oil price will remain 30% above your previous expectations”.

oil Brent). This guarantees comparability with the VAR literature in which most applications of impulse response functions are generated for a shock of one standard deviation.

3.4 Empirical results

We now discuss the results obtained from our survey-based impulse response analysis on the effects of oil price shocks. We structure the results section according to the level of aggregation. Starting with the aggregate results we compare the outcome of our survey experiment with a structural VAR (SVAR). In a next step, heterogeneity within our dataset is studied at the sectoral level. Thereafter, we study shock transmission channels at the industry level following Lee and Ni (2002). A further advantage of the survey experimental approach is illustrated at the firm level, where we dissect the generated survey based impulse responses by employing regression analysis. Thereby, we shed light on the relation between the oil consumption of firms and their responsiveness to oil price shocks.

3.4.1 Aggregate results

The definition of impulse responses generated from survey experiments is equivalent to a general definition of impulse responses in the time series literature (see Section 3.2). To compare survey-based impulse responses to time series impulse responses, we estimated a BVAR on oil prices, output (real GDP) and a producer price index (PPI). The VAR identification of oil price shocks follows Kilian (2008).²¹

Aggregate survey-based impulse responses are derived by aggregating the representative sample of firm-level data from our survey experiment using standard procedures (European Commission, 2007). First, individual impulse responses are aggregated to the industry level with firms' number of employees as weights.²² Second, a Swiss economy-wide aggregate is built from industry groups by utilizing gross value added shares as provided by the Swiss Federal Statistical Office. The aggregation scheme can be found in the appendix to this chapter at the end of the

²¹In contrast to Kilian (2008) we use Bayesian methods to estimate the model. See the appendix at the end of the thesis for further details. Note that our VAR based impulse responses depict average historical effects of oil price shocks whereas our survey based impulse responses determine the effects of a shock at the time of the survey in summer 2012, when the survey was conducted. This conceptual difference mutes comparability.

²²The number of employees is a proxy for the value added of a company.

thesis. By weighting the firm level data first for industry groups and second for the whole economy we ensure generalizability of our sample. We investigated whether our results vary with the applied weighting schemes and the results do not depend on the weighting scheme and unweighted results are only marginally different from weighted ones.

As can be seen from Table 3.1, the one standard deviation (= 30%) oil price shock has significant effects on the Swiss economy according to both approaches. The survey-based impulse response analysis yields a real turnover change of -0.4% within the first six months and by -0.6% within 18 months (see Table 3.3). The BVAR based impulse responses to an oil supply shock at their posterior mean range from -0.4% within the first 6 months to -2.8% within 18 months (see the appendix to this chapter at the end of the thesis for full fledged impulse responses). However, there is a lot of uncertainty around the BVAR estimates, with most of the probability mass in the negative region. BVAR error bands always include the values from the survey-based impulse responses.

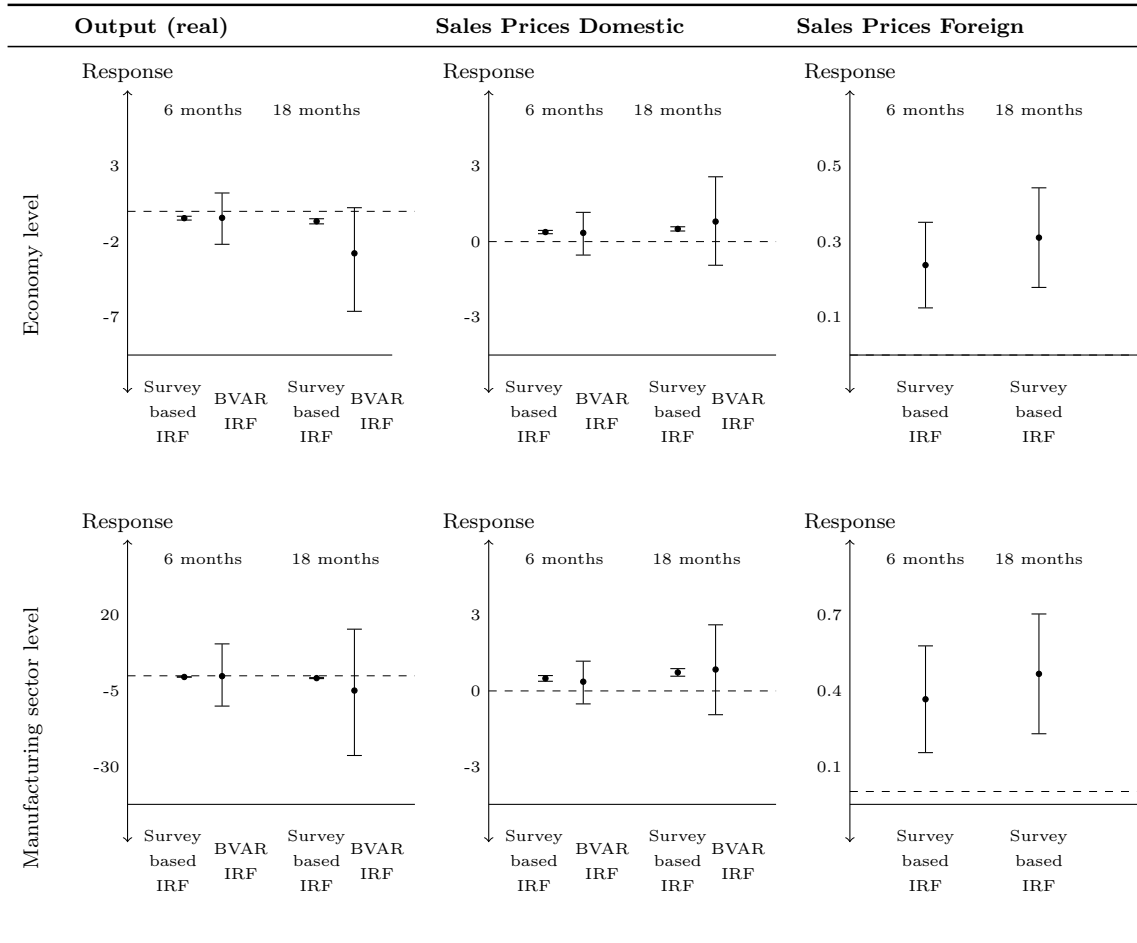
Domestic sale prices at the economy-wide level increase more strongly after 18 months (0.6%) than within the first 6 months (0.5%). These figures are strikingly in line with the responses obtained from the BVAR: price increases more strongly within 18 months (0.9%) than within 6 months (0.4%). Most of the posterior probability mass is in the positive region.

We also investigated the effects of an oil price shock on foreign sales prices. Foreign sales prices react weaker than domestic sales prices. Within 6 months foreign sales prices at the economy-wide level increase by roughly 0.3% , while we do not observe an additional increase within 18 months.

The economy-wide level survey-based impulse responses rely on turnover values, which are not the same as the GDP series used for the BVAR.²³ Thus, to provide further evidence for a comparison between VAR and survey-based impulse responses, we aggregate the survey experimental outcomes to the manufacturing sector level

²³A direct comparison between value-added GDP data and turnover data might be misleading (see, e.g., Kilian, 2008).

Table 3.1: Survey based and BVAR impulse response functions



The figure depicts impulse response functions (IRFs) derived from a survey experiment and IRFs generated by a BVAR in response to a one standard deviation oil price shock. Responses are displayed at the 6 and 18 months forecast horizon and are centered in between the whiskers. The whiskers of the survey based IRFs are ± 2 standard deviations. See the appendix to this chapter at the end of the thesis for calculation of these standard deviations. The whiskers of the BVAR impulse responses are the 5% and 95% percentiles of the respective posterior distribution of IRFs to an oil price shock.

and confront them with estimates from a BVAR using manufacturing real turnover and producer prices.²⁴

The manufacturing sector-level analysis confirms that the survey-based impulse responses and the VAR-based impulse responses are in the same ballpark. Real

²⁴The Swiss Federal Statistical Office only provides a manufacturing sector (but not an economy-wide) real turnover time series with sufficient length for VAR analysis.

turnover changes by -0.5% within 6 and -0.8% within 18 months according to the survey-based impulse response analysis. The BVAR yields posterior mean changes of 0.1% within 6 and -4.9% within 18 months. Note, however, that the BVAR manufacturing real turnover results are very imprecisely estimated, complicating the comparison of BVAR results and survey-based results. With respect to domestic sales prices, the values of the BVAR- and survey-based results are again in line with each other: Mean survey-based impulse responses depict a reaction of 0.8% to an oil price shock within six months and a response of 1.1% within 18 months. BVAR impulse responses lie at 0.4% within six months and 0.9% within 18 months. Survey-based impulse responses for foreign sales price reactions to an oil price shock on the manufacturing sector are stronger than on the economy wide level: Within six months foreign sales prices increase by 0.6% , within 18 months they rise by 0.7% .

We conclude that the impulse responses obtained from VAR analysis and the survey based impulse responses are broadly in line with each other.²⁵

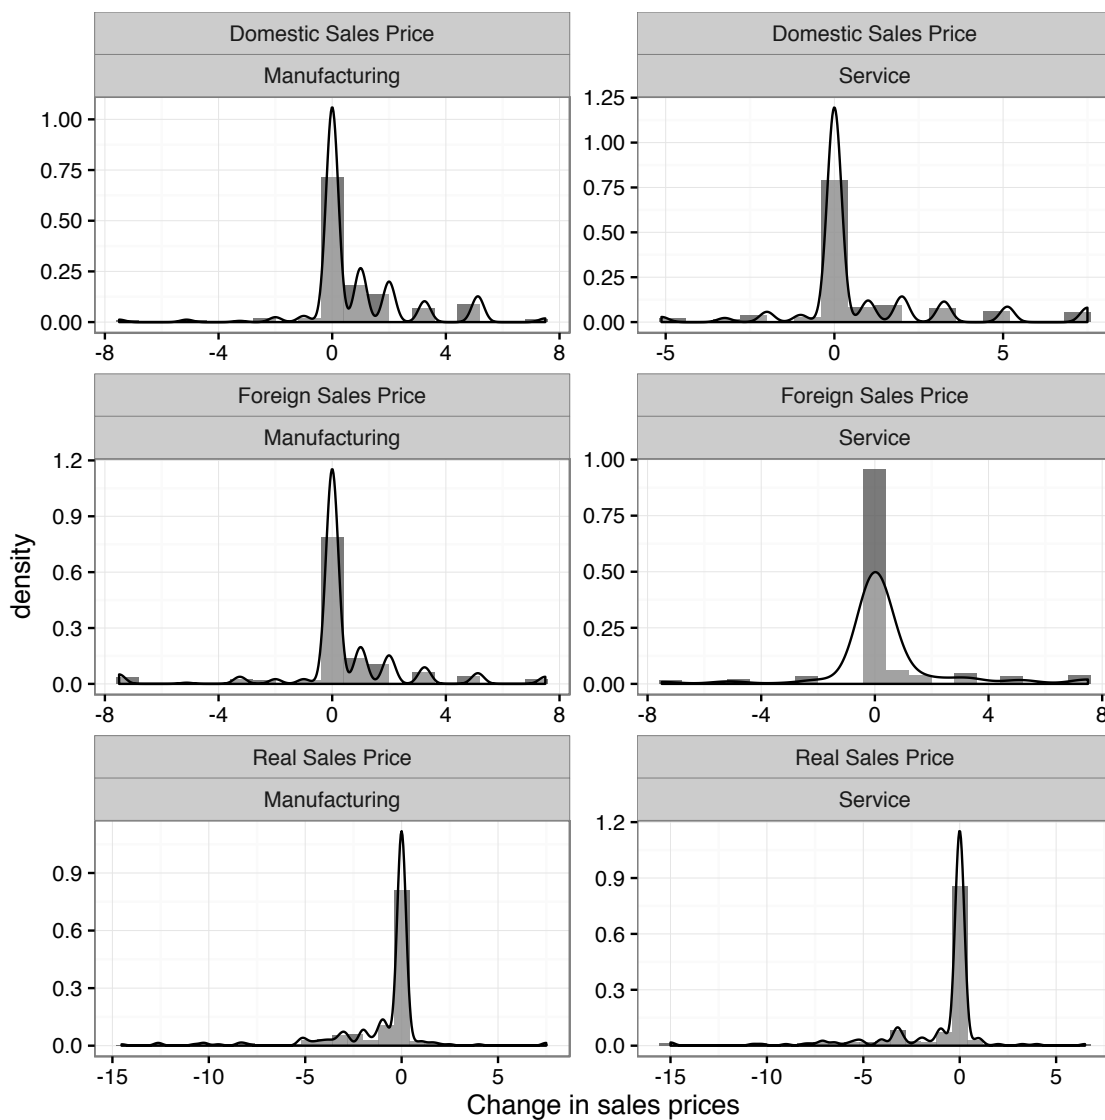
3.4.2 Heterogeneity within sectors

While aggregate results help characterize the response of an economy to macroeconomic shocks, a lot of heterogeneity might be hidden in aggregation. The results of our survey experiment shed light on the distribution of impulse responses on the firm level.

We observe a lot of heterogeneity within firms' impulse responses to an oil price shock. Figure 3.2 shows the empirical distributions of turnover changes in response to the oil price shock over all manufacturing sector firms and over all service sector firms. Distributions are presented in form of empirical probability mass functions (pmf) and in form of smoothed kernel densities calculated from the pmf. The overall weight of the real sales distributions for both the manufacturing and service sector lies in negative territory (upper row).

²⁵Survey-based impulse responses are also in line with the evidence in Peersman and Van Robays (2012). Among others, the authors estimated a SVAR for Switzerland based on the time period 1986-2010 and identified the effects of oil price shocks on real GDP and consumer prices.

Figure 3.2: Changes in real sales, foreign & domestic prices within 18 months



The histograms show 18 months impulse responses in real sales, foreign and domestic sales. The light bars show the relative frequency of firms' responses in percentage points. The transparent white areas represent smoothed kernel densities.

Both distributions have fat tails and are skewed to the left. The majority of firms only suffer from small reductions in output, yet a substantial fraction reports output losses of -5% and more in response to an unanticipated oil price shock of 30%. It is intriguing to see that large output reductions are more common for service sector firms than for manufacturing firms, while intuitively one might expect the opposite. The detailed analysis of firm responses on the industry level in Section 3.4.3 will shed light on this issue – the service sector firms most affected often belong to transportation & logistics.

Turning to foreign sales prices (middle row), most firms expect only slight changes in foreign prices compared to a no-shock scenario. Yet a gap becomes apparent between manufacturing firms and service sector firms. The right-hand tail of manufacturing firms' responses is fatter than for service sector firms. Also, the manufacturing sector distribution of impulse responses is obviously skewed to the right, while the service sector distribution is much less skewed. Furthermore, what distinguishes the pictures for foreign sales prices and real sales, is that there are quite a number of counterintuitive observations: a substantial number of firms reports *decreases* in foreign sales prices in response to an oil price shock. Section 3.4.3 covers this issue, revealing that some industries are affected via the foreign demand channel, which is responsible for decreasing sales prices in order to react to the declining demand for their products as the oil price increases.

The figures for domestic sales prices responses are broadly similar to the pictures for foreign sales prices. However, domestic sales price responses are more skewed to the right, both for manufacturing firms and for service sector firms. It appears, that firms are better able to raise sales volume in response to an oil price shock on the domestic markets than on foreign markets. Section 3.4.4 investigates what firm characteristics (such as oil dependency, market power, firm size, etc.) are responsible for price setting capabilities.

3.4.3 Industry level results

We now turn to the industry-level analysis of our survey experimental data. In theory, an oil supply shock affects output and prices through two main channels.

The cost-push channel is described by firms' reactions to a firm-specific cost increase in response to an oil price hike. The demand channel is characterized by reduced demand from firms' clients (consumers or other firms) as a response to an oil price increase. Following Lee and Ni (2002), we distinguish between the two channels in the following manner: After a negative oil supply shock the cost-push channel is important, where firms respond positively in terms of sales prices and negatively in terms of real sales volume. In contrast, the demand channel dominates when sales prices and real sales volume both react negatively.

Table 3.2 shows the pattern of impulse responses of output, domestic sales prices, and foreign sales prices (see Table 3.3 for more detailed figures). Industries are ranked according to their oil share in total expenditures. Similar to the SVAR setup of Lee and Ni (2002), we find that oil intensive industries, such as maintenance of machinery goods, chemicals & pharmaceuticals, and transport all show an increase in domestic sales prices and a decline in output, indicating the dominance of the domestic cost-push channel.²⁶ The two industries maintenance of machinery goods and chemicals & pharmaceuticals are dominated by the cost-push channel also on the foreign markets, where the sales price increase for chemical & pharmaceutical firms is by far the strongest of all industries on the foreign markets. Interestingly, the transportation & logistics industry, highly dependent on oil derivatives (e.g., fuel), is predominantly affected by a shift in international demand. This might be due to an elastic reaction of freight rates to an increase in oil prices, influencing the demand for international trade and logistics.

Four different patterns emerge among the remaining industries with medium or lower oil share. Industries with no consistent (i.e., insignificant) outcomes, those with a prevalence of the domestic cost-push channel, of domestic and foreign cost-push channels, and of foreign demand. Industries with a negligible shift in costs

²⁶In a VAR setup Herrera (2008) finds negative output effects of an oil price increase in highly oil intensive industries for U.S. data from 1959-2000. Industry level effects for the US have also been studied by Lagalo (2011). Jiménez-Rodríguez (2011) provide similar evidence for the UK, Germany, Spain, France, and Italy. In addition to the U.S. market Fukunaga et al. (2010) at industry level effects for Japan.

Table 3.2: Signs of impulse responses to an oil supply shock

| Industry | Peak effect on sales prices | Peak effect on real turnover | Dominating channels |
|---------------------------------|--------------------------------|---------------------------------|------------------------------|
| Maintenance of machinery goods | +/+ | – | Dom. & foreign cost-push |
| Chemicals & pharmaceuticals | +/+ | – | Dom. & foreign cost-push |
| Transport | +/- | – | Dom. cost-push & for. demand |
| Textiles | +/0 | 0 | Insignificant |
| Metals (except machinery) | +/0 | 0 | Insignificant |
| Construction | +/0 | 0 | Insignificant |
| Machinery & automobiles | 0/- | – | Foreign demand |
| Computers & electronics | +/0 | – | Dom. cost-push |
| Food & tobacco | +/+ | 0 | Insignificant |
| Wholesale & retail | +/+ | – | Dom. and foreign cost-push |
| Hotel & hospitality | 0/0 | – | Insignificant |
| Housing & technical services | 0/0 | 0 | Insignificant |
| Research & development | +/+ | – | Dom. and foreign cost-push |
| Telecom & IT services | +/+ | – | Dom. and foreign cost-push |
| Financial services & insurances | 0/0 | 0 | Insignificant |

Note: “0” denotes impulse responses that are not significantly different from zero. “+” means that the peak effect of the impulse response is significantly different from zero and positive. “–” means that the peak effect is significantly different from zero and negative. The left-hand side of the fraction “./.” refers to domestic sales prices and the right-hand side of the fraction to foreign sales prices. Peak effects are the largest reaction of impulses responses within 6 and 18 months.

or demand are textiles, metals, construction, food & tobacco, hotel & hospitality, housing & technical services, and banking & insurance.

Table 3.3: Industry level responses

| Industry | Turnover (real) | | Sales prices domestic | | Sales prices foreign | | Oil share |
|-----------------------------|-----------------|-----------------|-----------------------|----------------|----------------------|-----------------|----------------|
| | 6 months | 18 months | 6 months | 18 months | 6 months | 18 months | |
| Maintenance of machinery | -1.13 (0.37) | -1.67 (0.41) | 0.10 (0.10) | 0.60 (0.30) | 0.16 (0.16) | 0.55 (0.29) | 8.42 (2.72) |
| Chemicals & pharmaceuticals | -0.59 (0.28) | -1.13 (0.40) | 1.83 (0.29) | 2.24 (0.31) | 1.85 (0.39) | 2.20 (0.39) | 7.71 (1.18) |
| Transport | -0.55 (0.46) | -0.68 (0.46) | 0.77 (0.23) | 0.80 (0.23) | 0.23 (0.33) | 0.31 (0.31) | 3.54 (0.76) |
| Textiles | 0.38 (0.40) | -0.23 (0.33) | 0.60 (0.15) | 0.76 (0.18) | -0.08 (0.22) | 0.13 (0.18) | 3.33 (0.54) |
| Metals (except machinery) | -0.50 (0.47) | -0.73 (0.59) | 0.77 (0.33) | 1.03 (0.43) | 0.55 (0.30) | 0.41 (0.27) | 2.11 (0.47) |
| Construction | -0.93 (0.15) | -0.95 (0.17) | 0.96 (0.17) | 0.89 (0.22) | 0.14 (0.18) | 0.08 (0.21) | 2.08 (0.42) |
| Machinery & automobiles | -0.63 (0.23) | -0.63 (0.23) | 0.10 (0.17) | 0.22 (0.29) | -0.54 (0.47) | -0.73 (0.63) | 1.68 (0.37) |
| Computers & electronics | -0.06 (0.19) | -0.09 (0.24) | 0.58 (0.12) | 0.74 (0.14) | 0.11 (0.09) | 0.13 (0.10) | 2.56 (0.33) |
| Food & tobacco | -0.29 (0.13) | -0.54 (0.14) | -0.12 (0.14) | 0.13 (0.20) | -0.43 (0.25) | -0.16 (0.29) | 2.28 (0.32) |
| Wholesale & retail | -0.53 (0.21) | -0.82 (0.52) | 0.42 (0.13) | 0.54 (0.17) | 0.10 (0.19) | 0.16 (0.26) | 1.54 (0.32) |
| Hotels & hospitality | -0.36 (0.33) | -0.40 (0.34) | 0.06 (0.10) | 0.06 (0.12) | 0.06 (0.15) | 0.05 (0.16) | 1.58 (0.24) |
| Housing & techn. services | -0.76 (0.49) | -0.88 (0.50) | 0.18 (0.15) | 0.40 (0.19) | 0.20 (0.15) | 0.27 (0.16) | 1.73 (0.20) |
| Research & development | -0.64 (0.21) | -0.75 (0.26) | 0.61 (0.16) | 0.75 (0.17) | 0.33 (0.11) | 0.42 (0.14) | 1.56 (0.19) |
| Telecom & IT services | 0.05 (0.22) | -0.97 (0.48) | 0.00 (0.10) | 0.33 (0.18) | 0.00 (0.07) | 0.47 (0.24) | 1.14 (0.16) |
| Fin. services & insurances | -0.03 (0.04) | -0.03 (0.06) | 0.03 (0.03) | 0.03 (0.04) | 0.08 (0.09) | 0.08 (0.09) | 0.43 (0.11) |
| Manufacturing sector | -0.45 (0.12) | -0.82 (0.16) | 0.80 (0.09) | 1.05 (0.11) | 0.55 (0.12) | 0.73 (0.13) | 4.06 (0.35) |
| Service sector | -0.44 (0.09) | -0.58 (0.11) | 0.33 (0.06) | 0.42 (0.07) | 0.14 (0.06) | 0.20 (0.07) | 1.33 (0.10) |
| Economy total | -0.41 (0.07) | -0.61 (0.09) | 0.48 (0.05) | 0.62 (0.06) | 0.25 (0.06) | 0.34 (0.06) | 2.16 (0.14) |

The table depicts impulse responses of firms to an oil price shock in percent aggregated at industry, sector and economy level. Values in parenthesis are standard deviations (see appendix at the end of the thesis). Aggregation has been conducted as described in the text.

Interestingly, food & tobacco show an inelastic effect on output but are in the position to set prices domestically and internationally. While being able to set prices domestically, industries like textiles, metals, construction, and computers & electronic equipments are not able to set prices on international markets. Banks & insurances, having the lowest oil share in the entire sample, show no reaction in terms of output and prices. Industries where the domestic cost-push channel is predominantly active are manufactures of computers and electronic equipments. Domestic as well as foreign cost-push channels are important for wholesale & retail, research & development, and telecommunications & IT. In contrast, the foreign demand channel dominates for manufactures of machinery and automotive suppliers. This importance of the foreign demand channel in these sectors is due to the fact that Swiss automotive suppliers almost exclusively export their products to car producers in Germany, France, Italy, and the USA.

Industries with relatively high oil shares show comparatively strong output decreases in reaction to an oil price shock. However, table 3.3 reveals that there is no monotonic relationship between oil dependency and output reactions at the industry level. Lee and Ni (2002) also report no unambiguous link between oil dependency and the responsiveness of output. A reason might be that industry level results mix effects of oil dependency with industry-specific effects. Fortunately, the survey experimental approach allows us to zoom in further: in the next subsection, we show that by shifting the analysis to the firm level it is indeed possible to isolate oil dependency effects from industry-specific effects and to establish a clear link between oil dependency and output responsiveness.

3.4.4 Firm-level results

While common micro data inform about the heterogeneity of (unconditional) changes in output and prices, they lack insight about the variation in output and price changes in response to (i.e. conditional on) macroeconomic shocks. Survey experiments generate microeconomic data and, hence, combine the advantages of heterogeneity in micro data sets with structural identification of shock origins usually found in macro time series analysis. To demonstrate the usefulness of structural

microeconomic data generated by survey experiments we now turn to the dissection of driving forces behind impulse responses generated by survey experiments. Three questions and hypotheses are the focus of our analysis: i) While we find no clear relation between oil-price triggered output reactions and the oil dependency of industries (similar to the finding of Lee and Ni, 2002), we utilize our firm-level micro data set to investigate this relationship further and decompose oil dependency and industry class. ii) A firm's size and market power might play a role in explaining the magnitude of firm level responses to an oil price shock. Larger firms might possess a higher ability to absorb oil price shocks. Firms with more market power might be able to pass on cost increases to their customers more easily than firms with less market power. iii) Furthermore, sluggishness in responses might be expected, given that contracts are not fully flexible and rollover of contracts do not occur instantaneously.

Using the survey experimental data we can set up the following regression model:

$$\psi_{i,t+s} = \beta \mathbf{x}_i + \gamma \mathbf{z}_i + \theta \mathbf{d}_s + \xi_{i,s} \quad \forall i = 1, \dots, I. \quad (3.16)$$

where $\psi_{i,t+s}$ is either the real sales impulse response, domestic sales impulse response, foreign sales impulse response, purchase price impulse response or cost impulse response of firm i at horizon s to the 30% oil price shock (see the variables presented in this chapter's appendix at the end of the thesis). \mathbf{x}_i is a row vector of firm-specific explanatory variables and \mathbf{z}_i is a row vector of J industry dummy variables where the j -th = 1st, ..., J -th dummy variable takes value 1 if firm i is in industry j and zero otherwise. \mathbf{d}_s represents a time dummy which takes value 1 for $s = 18$ months and zero otherwise. \mathbf{x}_i includes three variables: firm i 's size as measured by its number of employees, firm i 's (non-shock scenario) oil input share defined as its expenditures on oil products (e.g., fuel, gasoline, diesel, oils, grease, plastics, chemical products) as a share of its total cost and firm i 's market power as measured by its profit margin, (i.e., (total sales – total costs)/total sales, on average over 2010–2012). β is a row vector of coefficients attached to \mathbf{x}_i , γ is a vector of industry-specific fixed effects that control for unobserved heterogeneity between industries, and θ captures the difference between impulse responses at the 18 month horizon and the 6 month

horizon. $\xi_{i,s}$ is the error term. And θ is a row vector of coefficients attached to \mathbf{D}_s . The regression coefficients are estimated via OLS.

As can be seen in Table 3.4, an increase in oil intensity by one percentage point intensifies the reduction of real sales in response to an oil price shock by 0.06 percentage points. For some industries such as textiles, chemicals & pharmaceuticals, transportation and retail, real sales are significantly reduced by between 0.9 and 1.3 percentage points (independent of oil dependency). However for other industries such as financial services & insurances or machine manufacturers we do not find significant effects on real sales which cannot be attributed to oil dependency. We find no significant time effects: real sales responses are statistically not different within 6 months and within 18 months after an oil price shock. Firm size and market power are insignificant in terms of real sales responses when controlling for industry effects.

Table 3.4: Response of real sales, domestic and foreign sales prices to oil price shock

| | Real sales | Domestic sales | Foreign sales | Purchase | Costs |
|---------------------------------|----------------------|---------------------|----------------------|---------------------|---------------------|
| | | prices | prices | prices | |
| | (1) | (2) | (3) | (4) | (5) |
| Oil input share | -0.061*** (0.018) | 0.077*** (0.012) | 0.014 (0.013) | 0.156*** (0.015) | 0.069*** (0.013) |
| No. of employees | 0.00003 (0.000) | -0.00001 (0.000) | -0.00002 (0.000) | 0.00000 (0.000) | -0.00001 (0.000) |
| Market power | -0.0002 (0.002) | 0.0001 (0.001) | 0.0001 (0.001) | 0.001 (0.001) | 0.001 (0.001) |
| Time dummy | -0.249 (0.171) | 0.213** (0.103) | 0.090 (0.128) | 0.322** (0.136) | 0.221* (0.117) |
| Food & tobacco | -0.027 (0.421) | 0.923*** (0.241) | 0.529* (0.309) | 0.951*** (0.317) | 1.512*** (0.289) |
| Textiles | -1.096*** (0.389) | 0.202 (0.236) | 0.101 (0.285) | 1.471*** (0.307) | 0.832*** (0.274) |
| Chemicals & pharmaceuticals | -0.928** (0.379) | 0.712*** (0.240) | 1.002*** (0.274) | 2.016*** (0.308) | 0.997*** (0.271) |
| Metals (except machinery) | -0.299 (0.294) | 0.222 (0.193) | 0.039 (0.225) | 0.865*** (0.258) | 0.727*** (0.224) |
| Computers & electronics | -0.516* (0.311) | 0.263 (0.200) | 0.001 (0.241) | 0.943*** (0.273) | 0.483** (0.235) |
| Machinery & automobiles | -0.092 (0.259) | -0.223 (0.176) | -0.099 (0.195) | 0.730*** (0.233) | 0.342* (0.200) |
| Maintenance of machinery goods | -0.293 (0.484) | -0.176 (0.322) | 0.299 (0.372) | 1.075** (0.443) | 0.400 (0.373) |
| Construction | -0.211 (0.355) | 0.476** (0.185) | 0.383 (0.273) | 0.941*** (0.244) | 0.758*** (0.213) |
| Whole sale & retail | -1.117*** (0.295) | 0.810*** (0.151) | 0.537** (0.224) | 0.976*** (0.197) | 0.783*** (0.170) |
| Transport | -1.302*** (0.364) | 0.969*** (0.202) | 0.929*** (0.259) | 1.822*** (0.264) | 1.052*** (0.226) |
| Hotel & hospitality | -0.131 (0.587) | -0.353 (0.308) | -1.345*** (0.446) | 1.647*** (0.395) | 0.037 (0.330) |
| Telecom & IT services | 0.237 (0.474) | -0.134 (0.278) | -0.299 (0.379) | 0.687* (0.372) | 0.849*** (0.319) |
| Financial services & insurances | -0.077 (0.289) | -0.117 (0.188) | 0.215 (0.218) | 0.101 (0.254) | 0.201 (0.217) |
| Housing & techn. services | -0.949*** (0.319) | 0.036 (0.190) | 0.150 (0.236) | -0.179 (0.257) | -0.260 (0.219) |
| Research & development | -0.407 (0.482) | -0.342 (0.288) | -0.121 (0.344) | 0.242 (0.375) | 0.139 (0.322) |
| Observations | 829 | 1,186 | 944 | 1,208 | 1,216 |
| Adjusted R ² | 0.133 | 0.196 | 0.058 | 0.407 | 0.216 |

p-values in parentheses. *** p<0.01, ** p<0.05, * p<0.1

We thus find that firms' responsiveness of output to an oil price shock clearly depends on the oil intensity of a firm, after controlling for covariates that also might influence firms' responsiveness. Our findings are in contrast to the ambiguous correlation between the oil share and industry-level responses reported in Lee and Ni (2002) and also found in the industry-level analysis in section 3.4.3. Moreover, we find no significantly stronger impulse responses after 18 month, reflected by the insignificant time dummy. The homogeneity of responses across time differs from the visual inspections of impulse responses for manufacturing firms in Lee and Ni (2002), where responses at a longer horizon are stronger than responses at a shorter horizon.

Turning to prices, we find a significant influence of the oil input share on domestic sales price responses. A one percentage point increase in the oil expenditure share increases the response of domestic sales prices to an oil price shock by 0.08 percentage points. In contrast, we do not find a statistically significant effect for foreign sales prices. Further, domestic sales prices react rather sluggishly to oil price shocks. The average domestic sales price increase from 6 months to 18 months after the shock is 0.21 percentage points, indicating a certain degree of stickiness in domestic sales prices. We further find differences in the price setting capabilities across industries. Firms belonging to industries such as chemicals & pharmaceuticals, construction, retail, and transportation seem to be able to set prices domestically independent of their oil share.²⁷ Moreover, we do not find any significant firm size or market power effects when controlling for industry effects.

Responses of purchase prices and costs induced by an oil price shock are also driven by firms' degree of oil dependency. A one percentage point increase in the oil share increases the purchase price response to a 30% oil price jump by 0.16% and the response of costs by 0.07%. Again firm size as expressed by the number of employees and market power do not have significant effects. A time effect is prevalent: the purchase price effect 18 month after the oil price shock is on average (after controlling for all other influencing factors) 0.32 percentage points higher than

²⁷By studying the frequency of price changes, Bils and Klenow (2004) and Nakamura and Steinsson (2008) also report heterogeneity in price setting capabilities.

the effect 6 months after the oil price shock. The costs response increases by 0.22 percentage points.

3.5 Conclusion

In this paper, we used a survey experiment to identify the effects of oil price shocks at the aggregate and firm level. We applied macroeconomic treatment scenarios to a representative sample of over 1000 Swiss firms. The prevalent diversity of industries and company types of our panel led to a rich dataset. The variations in firm responses allow drawing conclusions on the effects of aggregate shocks on the economy.

We showed that the results generated by our survey experimental procedure are conceptually equivalent to impulse responses calculated in VARs. However, in contrast to VARs, the impulse responses obtained from survey experiments allow for the identification of aggregate shocks without the need to impose parametric restrictions. Moreover, we showed that the identification assumptions in VARs imply restrictions on the firm-level impulse responses that might cause a bias at the aggregate level. This bias could potentially be tested using survey-based impulse response data.

At the aggregate level, we identified the impact of oil price shocks on Swiss economic activity and producer prices. We find that aggregate figures obtained from the survey experiment are in line with Bayesian VAR figures. We further decomposed the survey based impulse responses to an oil supply shock into a cost-push channel and a demand channel. It turns out that the impulse response channels are quite diverse across industries. At the firm level, we show a correlation between oil dependency and the responsiveness of output, that is not visible at the industry level, but which is identifiable when using our *structural* firm-level data.

We analyzed possible determinants of firm-level impulse responses by regressing them on a set of covariates such as oil intensity, market power, firm size, industry membership and time effects. Our findings suggest that market power and firm

size only play a minor role in explaining responses to an oil price shock, while oil dependency matters greatly.

Chapter 4

How Are Firms Affected by Exchange Rate Shocks? Evidence From Survey Based Impulse Responses¹

¹This chapter is based on KOF Working Paper 371.

4.1 Introduction

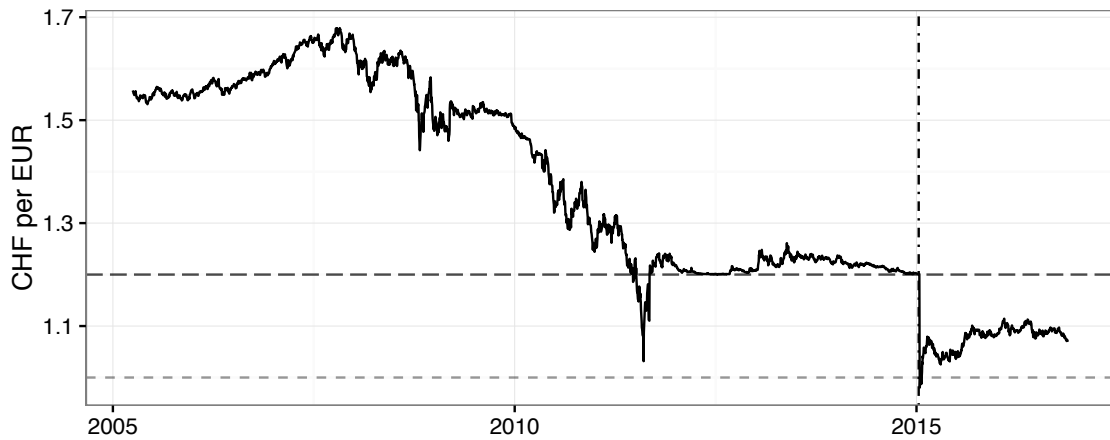
On January 15, 2015 the Swiss National Bank (SNB) announced the repeal of the Swiss Franc/ Euro exchange rate floor of 1.20 Francs per Euro. Immediately following the announcement the Swiss Franc appreciated significantly against the Euro and other currencies with high volatility. Two trading days later, the Swiss Franc/Euro exchange rate settled slightly above nominal parity. What are the consequences of such policy interventions and how does such a shock pass through to prices and profits of firms?

Figure 4.1 suggests the exchange rate shock induced by the SNB announcement is substantial: The swift appreciation of the Swiss Franc on January 15th, 2015 of almost 20% against the Euro is quite large in comparison with most exchange rate movements since the mid-1990s. Though the Swiss franc faced strong appreciation starting with the Great Recession in September 2008, which ended with the introduction of the 1.20 Francs per Euro floor in September 2011, it is not clear whether, or to what extent, this appreciation was unexpected and, hence, constitutes a macroeconomic shock.² The repeal of the exchange rate floor and the subsequent currency appreciation are likely to have considerable effects on the Swiss economy. However, given the peculiarity of the event and its presumably non-linear nature, it is difficult to derive effects based on historical time series. This paper proposes survey-based impulse response analysis as an alternative method to measure the effects of macroeconomic shocks. Further, we provide an application analyzing an exchange rate shock scenario.

The effects of exchange rate fluctuations on import prices and profits have so far been studied at the aggregate, as well as at the firm level. The primary objective of these studies is to gauge the size and speed of the adjustments to exchange rate movements. Most studies find an incomplete pass-through of exchange rates to prices. For instance, Campa and Goldberg (2005) perform a cross-country study of the effects of exchange rates on import prices. They present evidence for an

²It is also not clear whether the strong and sudden depreciation following the introduction of the exchange rate floor constitutes a macroeconomic shock.

Figure 4.1: CHF per EUR exchange rate



The graph depicts daily CHF per EUR exchange rates. The lighter, dashed, horizontal line indicates parity and the upper horizontal, dashed line marks the level of the exchange rate floor. The vertical line pinpoints the date of the repeal of the exchange rate floor. Source: quandl API (BOE).

incomplete pass-through in the short run and an significant higher one in the long run. Interestingly, their results are very heterogeneous across countries. Much lower effects are detected for the U.S. in the short run and the long run, whereas the pass-through for small open economy countries is significantly higher.

Using micro data for the U.S., Gopinath and Rigobon (2008) and Gopinath et al. (2010) deliver additional evidence on the transmission of exchange rate fluctuations. They find a decisive role of the currency choice of firms in the exchange rate pass through. The pass-through is very different for goods priced in dollars versus goods priced in other currencies. Both articles conclude that the pass-through is rather low for the U.S..

Lassmann (2013) provides empirical results for the correlation between exchange rate movements and imports prices and profits using time series survey data for Switzerland and she reports an increase in the probability of deteriorating profits, costs, and prices after an appreciation of the Swiss Franc. She also finds heterogeneous effects across firms, varying with the degree of international exposure.

The endogenous nature of exchange rates hampers the identification of causality and makes it difficult to specify an exchange rate shock that is orthogonal to any other economic innovation. VARs are well suited to analyzing aggregate multivariate time series and provide a framework for tackling potential endogeneity between different time series. Hahn (2003) and Faruqee (2004) apply structural VARs to aggregate Euro area data. They both find evidence for an incomplete pass-through of exchange rates on prices. Moreover, Hahn (2007) uses more disaggregated data to study the effects of exchange rate shocks for several sub-sectors and observes that the sub-sector with the strongest response, price-wise, is electricity, gas and water supply. By using a large scale macro model, Abrahamsen and Simmons-Süer (2011) analyze the effects of exchange rate fluctuations on many different macro variables. Among other things, their simulations imply a positive reaction of prices after a depreciation of the Swiss Franc. However, the size of this effect indicates a rather incomplete pass-through.

We contribute to existing literature in several ways. First, we provide firm-level dynamic causal effects, which we use to analyze the size and speed of the price adjustments in response to an exchange rate shock at the firm level. Survey-based impulse response analysis provides a convenient way to identify macroeconomic shocks. They create structural microeconomic data, which allows us to shed light on the heterogeneity of economic agents and allow to easily capture any kind of non-linearities. Survey-based impulse response analysis is on the spot, i.e., it determines the effects of shocks at the time when the survey was conducted rather than being based on historical time series. This feature makes the approach especially valuable in times of a structural break.

This a novel feature, given that we introduce the exogenous variations already at the firm level, retrieving the conditional and unconditional expectations of the firms directly. The method at hand produces dynamic causal effects (impulse responses) without imposing any identifying (parametric) restrictions. Further, a very general class of non-linearities across time and the cross-section can be handled easily. For example, we are able to analyze the heterogeneity of firms' responses to an exchange rate shock due to different export and import shares, firm sizes and industries.

Moreover, the firm-level approach also attenuates any omitted variable bias.³ More importantly, our approach allows us to dissect the causal effects of exchange rates on firms costs and profits using panel regression analysis, which allows us to draw inference on possible heterogeneity and non-linearity of the impulse responses to an exchange rate shock. Drechsel et al. (2015b) provide a more detailed discussion of the survey-based impulse response approach.

In July 2012, we conducted a survey-based impulse response analysis on the effects of a change in the Swiss Franc/Euro exchange rate floor as introduced by the Swiss National Bank in September 2011. We attached a questionnaire to the regular KOF Swiss Economic Institute Investment Survey, a major quantitative and qualitative statistical survey on firms' financial outcomes and investment plans in Switzerland. In a first step, firm representatives, mostly Chief Financial Officers (CFOs), Chief Executive Officers (CEOs) or heads of controlling, were reminded that the SNB had communicated that they would defend the exchange rate floor of 1.20 Francs per Euro and were asked to indicate their exchange rate expectations for the second half of 2012 and for 2013. In a second step, we asked firm representatives to evaluate the effects of – under otherwise unchanged macroeconomic circumstances – a change of the exchange rate floor from 1.20 francs to 1.10 francs per euro and a subsequent appreciation of the Swiss franc by the same magnitude on expected firm-specific turnovers and costs 6 months and 18 months ahead. Almost 900 firms completed the special questionnaire and our data is composed of representative samples of the populations of Swiss manufacturing, service and construction sector firms. In consequence, we can aggregate the resulting firm-specific impulse responses to higher levels via standard procedures for the construction of macroeconomic time series from representative micro data samples.

80% of all firm representatives forecasted the exchange rate floor of 1.20 francs per Euro to still be in place in 2013. More specifically, 60% of all firm representatives forecasted a Swiss Franc/Euro exchange rate of 1.20 in 2013, while 20% of them

³By conditioning on the actual information set of economic agents the scope of the information set is not an issue in survey-based impulse response analysis; selection of the right variables, under- or overidentification do not pose problems (see, e.g., Rudebusch, 1998).

forecasted an exchange rate depreciation up to 1.30 Swiss Francs per Euro. Interestingly, only 20% forecasted an appreciation. 11% of all firm representatives forecasted an exchange rate of 1.15 in 2013 and 9% of all firm representatives predicted an exchange rate of 1.10 in 2013. As firm representatives had been, immediately prior to the survey, explicitly reminded of the 1.20 Swiss Francs per Euro exchange rate floor and of the SNB's communicated willingness to defend this exchange rate floor, these forecasts are hardly the result of unawareness of the exchange rate floor. The only possible explanation is that the aforementioned firm representatives did, at the time of the survey in July 2012, indeed expected a repeal or a change of the exchange rate floor for 2013.

As far as firm representatives had forecasted the Swiss Francs/Euro exchange rate to be above 1.10, the scenario represents an exchange rate appreciation shock. We find that expected aggregate turnovers of the manufacturing sector decrease by 3.3% in the first six months after the exchange rate shock. After 18 months the expected decrease of aggregate manufacturing turnovers reaches 4.3% compared to the no shock state. The expected decrease in aggregate turnovers of the service sector is less severe, 1.6% after six months post-shock and 2.1% after 18 months. The construction sector is comparatively unaffected (0.4% and 0.8%). According to our findings, the exchange rate shock not only affects firms' turnovers but also firms' total costs, which is an approximate measure for import prices. Expected aggregate total costs of the manufacturing sector decline, compared to a no shock scenario, by 1.3% within six months and by 2.0% within 18 months. The respective expected aggregate total costs reductions for the service sector (construction sector) are 0.6% and 0.9% (0.2% and 0.5%). Furthermore, according to our findings the exchange rate shock leads to a substantial reduction in firms' profits. Expected total aggregate profits of the manufacturing sector decline by as much as 3.3% within six months due to the shock.⁴ The service sector experiences a reduction in expected total aggregate profits of 1.2%. The profit reduction in the construction sector is less severe (-0.2%).

⁴we did not ask for 18-months unconditional forecasts on turnover and costs, thereby we cannot compute unconditional profit expectations for 18 months and cannot derive expected profit responses within 18 months in response to an exchange rate shock.

The substantial negative profit effects, in particular in the manufacturing sector, are an indication of an incomplete pass-through which has also been observed by Campa and Goldberg (2005).

Notably, we find a lot of heterogeneity at the (sub-sector) industry level and at the firm level. For instance, the turnover of firms in the hotel industry (being part of the services sector) are strongly negatively affected by the exchange rate appreciation shock (despite the service sector aggregate effects being muted). Most of the turnover and cost distributions within a sector have fat tails and are skewed. A firm-level regression analysis reveals that a higher Euro area export share is associated with a stronger reduction in expected turnover in response to the exchange rate shock. Also, higher import shares from the Euro area, as well as higher import shares from the rest of the world, are associated with a stronger reduction in expected total costs in response to the exchange rate shock.

The remainder of the paper is structured as follows. The following section gives a brief review of the literature on exchange rate shocks. Section 4.2 describes the implementation of the survey-based impulse response approach. Section 4.3 presents the empirical results obtained from our survey-based impulse response analysis and Section 4.4 concludes.

4.2 The Survey-Based Impulse Response Approach

Following the previous chapter, which provides a more general discussion of the survey based impulse response approach, we create impulse responses from a survey that confronts firm executives with a scenario of a substantial shock in the CHF/EUR exchange rate. The scenario survey was conducted in July 2012 to assess the effects of a possible change in the exchange rate policy on the Swiss economy, as well as on specific sectors and at the firm level. Note that the survey was conducted when the exchange rate floor had already been in effect for several months.

The following subsection describes the implementation of the survey-based impulse response approach for the specific case of our exchange rate scenario. Section 4.2.1 describes the framework conditions of the survey itself, Section 4.2.2 explains

how our experimental design is incorporated into the questionnaire and will along with section 4.2.3 summarize the process of creating aggregate impulse response from the raw survey data.

4.2.1 Survey set-up

Our data was generated through a multi-mode survey which was attached to the summer 2012 wave of the semi-annual Investment Survey conducted by KOF Swiss Economic Institute. The KOF Investment survey is a major quantitative and qualitative firm-level survey on financial outcomes and planning of Swiss firms. Firms were able to answer either on paper or online using a LimeSurvey-based online questionnaire.⁵ The characteristics of the underlying sample are representative of the Swiss economy. Detailed information on the sampling procedure can be found in Appendix at the end of this dissertation.

893 Swiss firms completed the additional set of exchange rate related questions, of which 83 were from the construction sector, 398 from manufacturing sector and 412 from service sector. Our data constitute representative samples of the populations of Swiss construction, manufacturing and service sector firms. Nevertheless the data do not provide a representative sample of the total population of Swiss firms. The aforementioned sectors cover only 60% – 70% of economy-wide value added; however, other sectors (e.g. private households, public administration, social services and health care) are not included in our survey. Firms' responses come mostly from CFOs, CEOs, and heads of controlling. Respondents are taking part in KOF enterprise surveys on a regular basis and are accustomed to KOF questionnaire design. In order to ensure the relevance of our questions to practitioners, we conducted an interviewer pre-test with a group of selected firms.

Participating firms received an invitation letter and the questionnaire on paper and in electronic format in order to facilitate their participation. Anonymity of responses has been guaranteed and all KOF surveys are subject to Swiss statistics

⁵LimeSurvey (www.limesurvey.org) is a free open source software that can be used to host and run online surveys.

law. If addressed participants did not respond within 2.5 weeks they received a reminder. Firms received a second reminder via telephone two weeks later if they still did not respond. Questionnaires were sent out in German, French and Italian.

4.2.2 Experimental design

Before participants answered the scenario questions, participants were asked to provide details on their firm's key financial figures for 2010, 2011 and the 1st half of 2012, as well as expected figures for the 2nd half of 2012 and 2013. This step is helpful to set a benchmark for the scenario analysis. Also, representatives were asked to provide exchange rate forecasts before being exposed to scenario (unconditional). We reminded firms of the exchange rate floor of 1.20 CHF per Euro which was in effect at the time of the survey. Note that we refer to expectations that are not subject to our scenario as *unconditional*.

In a next step, the questionnaire confronted firm representatives with the counterfactual scenario. Firms' representatives were asked to imagine an exchange rate innovation, namely – ceteris paribus – a change in the exchange rate floor from 1.20 to 1.10 Swiss Franc per Euro and a subsequent appreciation of the Swiss franc. Firm representatives were then asked to evaluate the effect of the exchange rate innovation on their firms' key financial figures:

Suppose, the SNB shifts the Swiss franc/euro exchange rate floor to 1.10 Swiss francs per euro under otherwise constant economic circumstances. As a consequence the exchange rate moves to 1.10 Swiss francs per euro, which is an appreciation of the Swiss franc. Please indicate how your financial figures change compared to your previous expectations for these figures.

As an excerpt of the questionnaire, the answer options for total turnover are as follows (the complete questionnaire can be found in the appendix to chapter 3 at the of this dissertation.):

| <u>2nd Term 2012</u> | | | | | | | | | | | |
|---------------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|-------------------------------------|
| ≤-7.5% | -5% | -3% | -2% | -1% | 0% | 1% | 2% | 3% | 5% | ≥7.5% | N/A |
| <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> |

| <u>2013</u> | | | | | | | | | | | |
|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|-------------------------------------|
| ≤-7.5% | -5% | -3% | -2% | -1% | 0% | 1% | 2% | 3% | 5% | ≥7.5% | N/A |
| <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> |

Before assume that the above scenario means an appreciation shock for Swiss firms we need to consider the following argument: If a representative unconditionally expects the exchange rate to be at 1.10, the shock would just meet their expectations and would thus not be considered a shock. The ability to absorb this exchange rate innovation would still depend on the firm's market power and hedging ability. However, as section 4.3.1 explains in greater details, 90 % of the firms indeed expected the exchange rate to stay well above 1.10 CHF per EUR. In other words, the unconditional exchange rate expectations provide empirical evidence that most firms would consider the scenario to be an appreciation shock.

The scenario question is designed such that answers represent the firm-specific dynamic causal effects, or treatment effects that this exchange rate shock would have. This effect can be expressed formally as

$$\delta_{i,s} = E(y_{i,t+s}|\eta_{i,t} = 1) - E(y_{i,t+s}|\eta_{i,t} = 0), \quad \text{for } s = 6 \text{ months, } 18 \text{ months,}$$

where $\eta_{i,t}$ is the treatment variable with $\eta_{i,t} = 1$ when firm i receives the shock treatment at time t (treatment scenario) and $\eta_{i,t} = 0$ when the firm does not receive the shock treatment at time t (control scenario). Further, at time t , $E_t(y_{i,t+s}|\eta_{i,t} = 1)$ is firm i 's expected total turnover for horizon s given that the exchange rate shock occurred at time t and $E_t(y_{i,t+s}|\eta_{i,t} = 0)$ is its, at time t , expected total turnover for horizon s given the exchange rate shock did not occur ceteris paribus. The dynamic causal effect described in the above equation is equivalent to the definition of impulse responses that is prevalent in the time series literature, where the treatment is an unanticipated (aggregate) shock at time t with its effects s periods after the shock has occurred (see, e.g. Hamilton, 1994). Thus, we refer to $\delta_{i,s}$ as the *survey-based impulse response*.

In the same manner the questionnaire asked firm representatives to estimate the effect of the exchange rate shock on their total cost for the second half of 2012 (within 6 months) and for 2013 (within 18 months). The questionnaire also asked for firms' exports in terms of total turnover ("export share") and imports in terms of total expenses ("import share"), to/from the Euro area and to/from the rest of the world. The appendix to this chapter at the end of the dissertation gives a comprehensive list of variables used in this study.

By asking the *same* firm representative about a control scenario and a treatment scenario, the survey impulse response analysis follows a within-subject design (see Charness et al., 2012). One might also follow a between-subjects design by randomly assigning different scenarios to firms. Pre-tests yielded that it is more convenient for firms to indicate the change in projections under the treatment scenario compared to the control scenario, rather than indicating projections under the treatment scenario and once again under the control scenario.

The question that remains is: Can firm representatives know how projections would change if their firm was hit by an exchange rate shock? As argued by, among others, Gaines et al. (2006) the treatment scenario must be realistic in the sense that respondents have been confronted with similar scenarios in the past and/or that they have previously considered the scenario and its effects before. We consider this to be the case since firm executives have had to cope with recurring swift and partly unanticipated, exchange rate movements in the past. The appendix to this chapter at the end of the thesis provides a discussion about further validity issues.

4.2.3 From firm-level to aggregate impulse responses

The data generated from our firm-level impulse response analysis constitute representative samples of the population of Swiss construction, manufacturing and service sector firms. Aggregate survey-based impulse responses are derived via standard procedures (see European Commission, 2007). First, individual impulse responses are aggregated to the industry level (based on NACE sector classifications) with firms' number of employees as weights.

Second, aggregate impulse responses for the manufacturing, service and construction sector are built from sub-sector industry groups using gross value added shares provided by the Swiss Federal Statistical Office.⁶ We refrain from building economy-wide aggregate impulse responses since the aforementioned sectors cover only 60% – 70% of total Swiss value added and other sectors are not included in our survey (see Section 4.2.1). The aggregation scheme is relegated to the appendix to this chapter at the end of the thesis.

As a robustness check, we investigated the dependency of our sector-level aggregate results in relation to the applied weighting schemes. The sector-level aggregate results are robust to different weighting schemes and unweighted results are only marginally different from weighted ones.

4.3 Results

This section presents the empirical results obtained from our survey-based impulse response analysis. The upcoming subsections structure the different aspects of these results in the following way: Section 4.3.1 presents firms' exchange rate expectations in a first preliminary step to contextualize our findings with firms' mindset back in July 2012. Section 4.3.2 exploits one of the main benefits of the survey based impulse response approach and presents results at various aggregation levels, namely the industry and sector level. Section 4.3.3 breaks our findings down even further discussing firm-level results, including a firm-level panel regression.

Recall from Section 4.2 that we study the effects of exposing firms to the following shock: an unexpected change of the Swiss franc/Euro exchange rate floor from 1.20 to 1.10 CHF / Euro all else being equal, and a subsequent appreciation of the Swiss franc from 1.20 to 1.10 CHF / Euro.

⁶See Bundesamt für Statistik, Produktionskonto nach Branchen, Table 3a.3.

4.3.1 Exchange rate expectations

Before setting up the actual scenario of an exchange rate shock we asked firms about their current exchange rate expectations. We also reminded firms of the existing exchange rate floor and the SNB's announcement to defend the exchange rate floor of 1.2 Swiss franc per Euro. In other words, we tested firms' trust in the announcement of the SNB and thus also whether they would be hit by a shock when the exchange rate floor was changed or repealed. For the second half of 2012, the overall majority of firm representatives expected the exchange rate to stay at or closely above 1.20 Swiss francs per Euro.


For 2013, 80% of all respondents still expected the exchange rate floor of 1.20 Swiss francs per Euro to be in effect. Figure 4.2 depicts the distribution of the (unconditional) exchange rate expectations for 2013. To be more specific, 60% of all firm respondents forecasted a Swiss franc/Euro exchange rate of 1.20 in 2013, while 20% forecasted an exchange rate depreciation up to 1.30 Swiss francs per Euro. Strikingly, 11% of all firm respondents forecasted an exchange rate of 1.15 in 2013 and 9% of all firm respondents predicted an exchange rate of 1.10 in 2013. Although this group of firms, which anticipated an appreciation is an interesting subgroup for further research, our discussion will mostly focus on the lion's share of companies which did not anticipate an appreciation shock and were thus exposed to a shock in our experiment.

4.3.2 Sector and Industry Level Effects

This section discusses firms' turnover, cost and profit reactions to the exchange rate shock given in our scenario. Turnovers and costs decline in all sectors. Depending on firms' export and import shares, turnovers and costs decline at different rates across sectors. The effect on profit is not clear, a priori, because lowered import costs may outweigh lower turnovers, or vice versa. Thus, we will first discuss cost and turnovers and turn to profits later in this section.

Note that the upcoming findings do not mean that a specific sector expects an overall decline in turnovers or costs. Instead, the figures indicate a reduction

Figure 4.2: Exchange rate expectations within 18 months



The bar chart shows the relative share of firms' expectations with respect to the average CHF / Euro exchange rate within 18 months. The X-axis displays exchange rate CHF / Euro exchange rates and the Y-axis depicts the share of firms' answers.

in turnover and costs compared to the turnovers and costs that would have been expected in case of the no shock scenario (= exchange rate floor stays at 1.20 Swiss francs per Euro). An example: If turnover increases by 2% within six months in case of no shock, a decline in response to the specified exchange rate shock, of 1% within six months leads to an overall turnover growth of $1.02 \cdot 0.99 - 1 = 0.98\%$.

Manufacturing turnover is expected to decrease by 3.3% in the first six months after the exchange rate shock compared to the no shock scenario.⁷ After 18 months the expected decrease of aggregate manufacturing turnover reaches 4.3% compared to a state of the world in which no shock has occurred. The service sector's expected decrease, six months after the shock, is 1.6%. After 18 months, it reaches 2.1%. The construction sector is less affected, within six months expected turnover only

⁷See also table 4.1 for a comprehensive overview.


declines by 0.4% and 0.8% within 18 months. The manufacturing sector strongly depends on exports to the Euro area and its (aggregate) export share is 36.3% in the sample. With a sample export share of 7.2%, the service sector depends less on exports to the Euro area, and the construction sector with a sample export share to the Euro area of only 2.2% depends the least on exports to the EU. All described weighted averages are statistically different from zero at the 1% significance level.

An appreciation of the Swiss franc not only causes lower turnovers, compared to a no-shock scenario, import costs and thereby total costs, will be reduced. For the manufacturing sector, total costs are expected to decline, compared to a no shock scenario, within six months and 18 months after the exchange rate shock. The expected decline for the 18-months ahead period (2.0%) is stronger than for six months ahead (1.3%). Both effects are statistically different from zero at the 1% significance level. Given an (aggregate) sample import share from the Euro area of 30.1% an appreciation of the Swiss franc helps manufacturing firms to save on costs. The service sector and the construction sector depend less on imports from the Euro area with (aggregate) sample import shares of 14.5% and 6.2%, respectively. In consequence, there is less room to reduce costs in response to the exchange rate shock in these sectors. For the service sector, expected total costs decrease by 0.6% within six months and by 0.9% within 18 months, while the construction sector's expected costs decline by 0.2% within six months and by 0.5% within 18 months. With the exception of the 6-month effect for the construction sector, all responses are statistically different from zero at the 1% significance level.

The heterogeneous turnover and cost reactions can also be seen within sectors. Figure 4.3 and 4.4 show the correlations between firms' turnover and export shares and correlations between firms' costs and import shares, respectively. Without controlling for any covariate, we find a clear negative relationship and substantial differences in the degree to which different industries are affected.

A sector-by-sector break down of these industry-level effects illustrates the benefits of the survey-based impulse approach for applied research. The service sector for example has only shown limited reactions to the exchange rate shock scenario on the aggregated level. Still, the hotel industry as part of the service sector, ranks


Figure 4.3: Turnover 18 months vs. export shares



The figure plots the correlation of expected turnover and sectors' export shares to the Euro area for a 18 month ahead horizon. The dots represent single sectors while the dotted line shows the general correlation.

among the Swiss industries most strongly affected (-2.9% decline in turnover for the 6-month horizon and -3.8% for the 18th-months horizon respectively) by our shock scenario. An appreciation of the Swiss franc affects not only tourism with its dependency on foreign tourists, but also gives Swiss tourists an incentive to go abroad. Though tourism is not among the top industries in the country in terms

Figure 4.4: Costs 18 months vs. import shares



The figure plots the correlation of expected costs and sectors' import shares from the Euro area for a 18 month ahead horizon. The dots represent single sectors while the dotted line shows the general correlation.

of value added, tourism is very important for particular regions and contributes to Switzerland's world-reknoned image, justifying a disaggregate look.⁸

⁸The reported export share seems very low for the tourism industry and is caused by unusual item non-response and zero answers in the export share question. The high item non-response is likely caused by the general wording of the term export share for all participants. The term 'export share' rather than 'share of foreign guests' likely confuses participants from the tourism industry and leads to item non-response and dubious answers. Similar confusion is observed in other survey projects when using the concept of 'capacity utilization' instead of the more specific 'occupancy rates'. Hence results for tourism with respect to export shares should not be taken too seriously. However, there is no indication that the impulse responses of the tourism industry are affected by this confusion.

Table 4.1 provides a comprehensive overview of cost, turnover and profit effects. The heterogeneity in firms' responses can be observed also within the manufacturing sector. The industries reacting the strongest in terms of expected total turnover are machine producers & car suppliers (-4.0% within six months, -5.3% within 18 months), as well as metal producers (-4.2% within six months, -5.2% within 18 months). This is not surprising given that Euro area export shares in the sample are also quite large, with 37% for machine producers & car suppliers and 40% for metal producers. Strong effects can also be observed for chemical & pharmaceutical firms (sample export share Euro area of 43%), firms in the electro-technical & fine mechanics industry (sample export share Euro area: 36%), as well as firms belonging to the textile industry and firms maintaining machinery goods (sample export share Euro area of 30% and 27%). The negative expected turnover effects compared to a no shock scenario for chemical & pharmaceutical firms amount to 3.3% within six months and 3.7% within 18 months. Most of the negative effect in chemicals & pharmaceuticals is due to a large negative effect for chemical firms, while pharmaceutical firms only report minor losses in total expected turnover. The expected turnover reduction for the electro-technical & fine mechanical firms amounts to 2.7% within six months and 4.7% within 18 months. The expected turnover effects for textile firms and firms maintaining machinery goods reach -4.1% and -4.0% within six months and -4.9% and -5.3% within 18 months.

Those firms expecting to suffer strongly from the exchange rate shock in terms of total turnover often also expect to benefit substantially from the reduction in total costs. Machine producers & car suppliers, with a sample import share from the Euro area of 38% , expect considerable cost reductions, namely 1.9% within six months and 2.6% within 18 months. Metal producers' expected costs (sample import share Euro area of 39%) decline by 1.9% within six months and by 2.3% within 18 months. Chemical & pharmaceutical firms' expected costs (import share Euro area of 30%) reduce by 0.7% within six months and by 0.9% within 18 months. Further, the expected costs of the electro-technical & fine mechanical firms (sample import share Euro area of 24.8%) decrease by 1.2% within six months and by 3.0% within 18 months. Firms belonging to the textile industry and firms maintaining machinery

goods (sample import shares Euro area of 29.6% and 21.9%) expect to benefit from reduced costs by 1.5% and 2.2% within six months and by 2.1% and 2.3% within 18 months. As stated above, all figures give the change in expected costs due the shock as compared to the no shock scenario.

Table 4.1: Industry and sector level survey based impulse responses

| Industry/sector | Turnover | | Costs | | Profits | Euro area share | | CHF/EUR |
|------------------------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|
| | 6 months | 18 months | 6 months | 18 months | 6 months | Exports | Imports | 18 months |
| Food & tobacco | -1.96 (0.45) | -2.00 (0.42) | -0.95 (0.36) | -1.65 (0.51) | -2.77 (0.49) | 24.62 (3.39) | 24.01 (3.17) | 1.20 (0.007) |
| Textiles | -4.10 (0.52) | -4.86 (0.55) | -1.53 (0.44) | -2.07 (0.53) | -3.83 (0.58) | 30.01 (4.03) | 37.15 (4.18) | 1.19 (0.007) |
| Chemicals & pharma | -3.29 (0.43) | -3.67 (0.47) | -0.74 (0.27) | -0.93 (0.31) | -3.05 (0.43) | 43.32 (2.76) | 29.59 (3.01) | 1.20 (0.002) |
| Metals (except machinery) | -4.19 (0.30) | -5.20 (0.31) | -1.85 (0.33) | -2.30 (0.40) | -3.72 (0.36) | 39.55 (3.02) | 39.18 (3.15) | 1.21 (0.007) |
| Electronics & fine mechanics | -2.73 (0.39) | -4.67 (0.30) | -1.16 (0.32) | -2.95 (0.38) | -2.69 (0.46) | 36.14 (2.95) | 24.82 (2.46) | 1.20 (0.006) |
| Machinery & automobiles | -4.02 (0.27) | -5.29 (0.28) | -1.91 (0.26) | -2.60 (0.34) | -4.18 (0.31) | 36.98 (2.18) | 37.93 (2.18) | 1.20 (0.006) |
| Repair of machineries | -4.13 (0.46) | -4.75 (0.41) | -2.19 (0.45) | -2.33 (0.47) | -4.36 (0.57) | 26.99 (4.21) | 21.86 (2.13) | 1.19 (0.010) |
| <i>Manufacturing sector</i> | -3.30 (0.16) | -4.25 (0.16) | -1.27 (0.13) | -2.04 (0.16) | -3.27 (0.17) | 36.32 (1.2) | 30.11 (1.17) | 1.20 (0.002) |
| Wholesale & retail | -2.09 (0.22) | -2.66 (0.25) | -0.68 (0.15) | -1.25 (0.17) | -1.86 (0.27) | 6.23 (1.28) | 27.67 (2.48) | 1.20 (0.003) |
| Transport | -1.31 (0.22) | -1.46 (0.24) | -0.96 (0.19) | -1.08 (0.2) | -1.35 (0.24) | 10.08 (1.69) | 7.58 (2.61) | 1.18 (0.009) |
| Hotels & restaurants | -2.88 (0.51) | -3.82 (0.66) | -0.81 (0.29) | -1.26 (0.36) | -2.81 (0.56) | 1.60 (1.51) | 7.92 (2.02) | 1.22 (0.013) |
| Media & IT services | -1.07 (0.15) | -1.13 (0.16) | -1.90 (0.38) | -1.43 (0.33) | -0.85 (0.34) | 7.66 (1.20) | 19.86 (3.69) | 1.20 (0.005) |
| Banking & insurance | -1.00 (0.22) | -1.19 (0.26) | -0.34 (0.19) | -0.31 (0.21) | -0.59 (0.24) | 3.31 (1.28) | 2.56 (0.66) | 1.20 (0.003) |
| Housing & tech. services | -1.25 (0.23) | -2.44 (0.31) | -0.28 (0.20) | -0.32 (0.19) | -0.74 (0.25) | 16.82 (2.47) | 7.10 (2.14) | 1.20 (0.003) |
| Research & development | -2.06 (0.38) | -3.37 (0.53) | 0.05 (0.54) | -0.38 (0.68) | -0.33 (0.68) | 6.22 (2.96) | 8.83 (2.65) | 1.13 (0.012) |
| <i>Service sector</i> | -1.60 (0.11) | -2.14 (0.13) | -0.64 (0.09) | -0.85 (0.10) | -1.24 (0.13) | 7.19 (0.72) | 14.48 (1.15) | 1.19 (0.002) |
| <i>Construction sector</i> | -0.41 (0.18) | -0.84 (0.23) | -0.15 (0.09) | -0.52 (0.18) | -0.23 (0.15) | 2.15 (0.55) | 6.23 (0.92) | 1.20 (0.006) |

The table depicts survey based impulse responses of firms to an exchange rate shock in percent aggregated at industry or sector level. Values in parenthesis are standard deviations (see Appendix to this chapter at the end of the thesis.). Aggregation has been conducted as described in Section 4.2.

As stated before, the effect on profits is not clear a priori. If turnover effects outweigh cost effects a reduction in profits might be expected and vice versa. We derive profits for the within-6-months period as the absolute value of profits to be expected after the exchange rate shock within 6 months divided by unconditionally (before the exchange rate shock) expected absolute profits within 6 months (i.e., until by the end of 2012).

$$\Delta\Pi_{i,s=6} = \frac{Y_{i,s=6}^{unc} \times \Delta Y_{i,s=6}^c - C_{i,s=6}^{unc} \times \Delta C_{i,s=6}^c}{Y_{i,s=6}^{unc} - C_{i,s=6}^{unc}},$$

with firms $i = 1, \dots, I$. $Y_{i,s=6}^{unc}$ and $C_{i,s=6}^{unc}$ are the unconditionally expected absolute turnover (Y) or costs (C) of firm i at horizon $s = 6$ months. $\Delta Y_{i,s=6}^c$ and $\Delta C_{i,s=6}^c$ are the survey-based impulse response in either turnover or cost of firm i at horizon $s = 6$ months, conditional on the exchange rate shock. Superscript c indicates conditional forecasts (on shock occurrence), superscript unc unconditional forecasts (i.e., no shock occurs). $\Delta\Pi_{i,s=6}$ is the growth rate of profits Π for firm i at horizon $s = 6$ months (i.e., by the end of 2012) conditional on the exchange rate shock.

It turns out that negative turnover effects outweigh the reductions in cost within 6 months at the sectoral level. Profits of manufacturing firms decline by -3.3%, while service sector firms' profits decrease by -1.24%. The construction sector is not affected in terms of profit, the aggregated value is not statistically different from zero.

A similar result can be found at the more finely grained industry level. The results in Table 4.1 show that for every subgroup within the manufacturing sector and for every industry in the service sector the negative turnover effects outweigh reductions in costs. All industries in the manufacturing sector expect substantial reductions in profits within 6 months. Specifically, profits are expected to decline by 4.2% for machine producers & car suppliers, by 4.4% for firms repairing machinery goods, by 3.8% in the textile industry, by 3.7% for metal processing, by 3.1% for chemical & pharmaceutical firms and by 2.8% and 2.7% for food & tobacco as well as electronics and fine mechanics. While all industries in the service sector expect a reduction in profits within 6 months, the effects are less pronounced than for the manufacturing sector. The only exception are hotels and restaurants. Their

profits are expected to decline by 2.8% within 6 months. Wholesale trades and retailers expect profits to decrease by -1.9% within 6 months and transportation and logistic companies expect a decline by -1.4%. All other service sector industries (media, telecommunication & IT; banking, financial services & insurance; housing, freelancing, advisory, architects, tech. services) expect reduced profits between 0.6% and 0.9%. Lastly, the effect for research and development is statistically not different from zero.

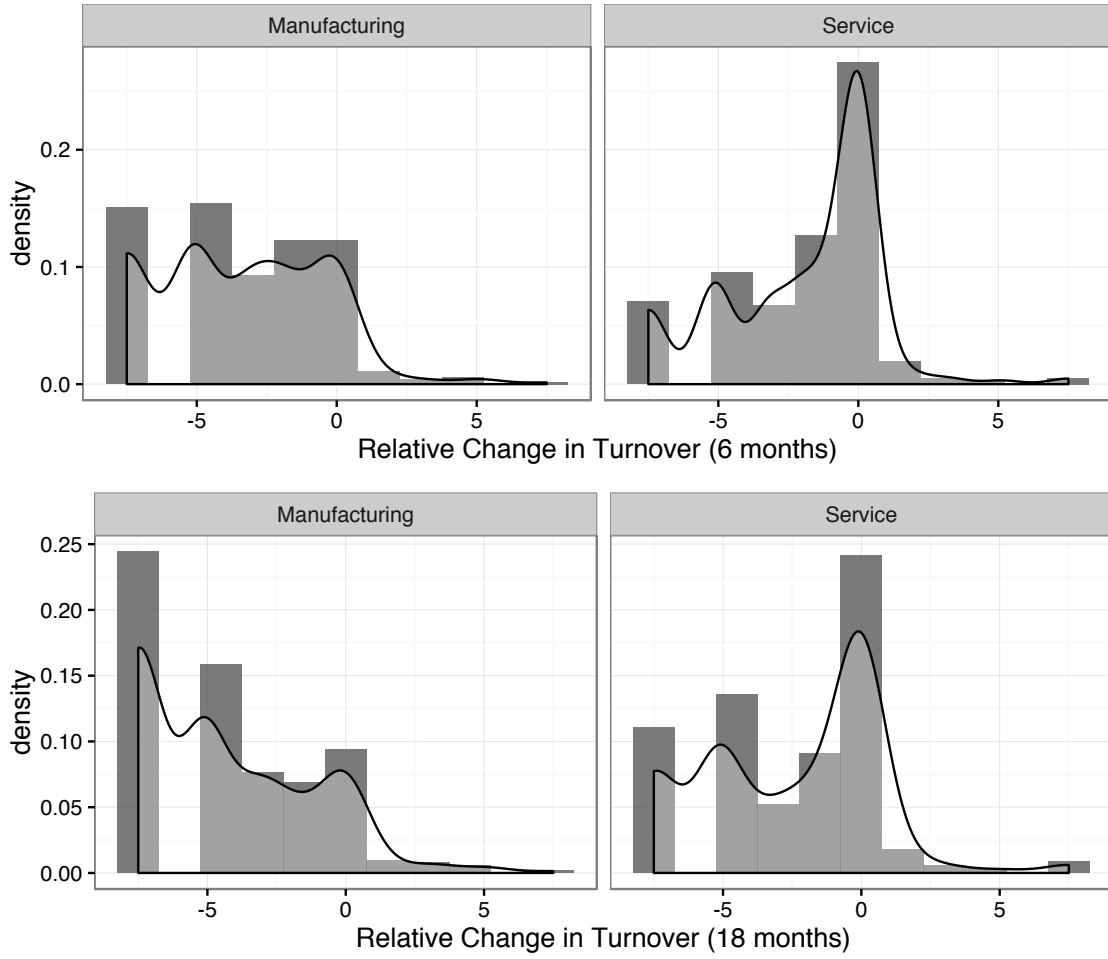
4.3.3 Firm-level heterogeneity

We observe a lot of firm-level heterogeneity in the collected data set. Figure 4.5 shows the empirical distributions of turnover changes in response to the exchange rate shock over all manufacturing and service sector firms. Distributions are presented in form of empirical probability mass functions (pmf) and smoothed kernel densities calculated from the pmf. While the overall weights of the distributions lie in the negative range for both sectors, the shift is more pronounced for the manufacturing sector. All distributions have fat tails and are skewed to the right. While the 6- and 18-months distributions for service sector firms might still be classified as skewed normal distributions with fat tails, the 6- and 18-months distributions for manufacturing sector firms are far removed from normal distributions. As can be seen from Figure 4.6, cost impulse responses also show substantial dispersion. However, the distributions are less skewed than for turnover responses.

4.3.4 Firm-level regression analysis

In order to investigate the driving forces behind the survey-based impulse responses of firms to the exchange rate shock, we conduct panel regression analysis. Controlling for industry and previously expressed exchange rate expectations, we find a significant effect of the Euro area export share on expected turnover impulse responses to the exchange rate shock. Furthermore, we find significant effects for import shares and exchange rate expectations on cost impulse responses and significant effects for trade surpluses with the Euro area.

Figure 4.5: Turnover distributions



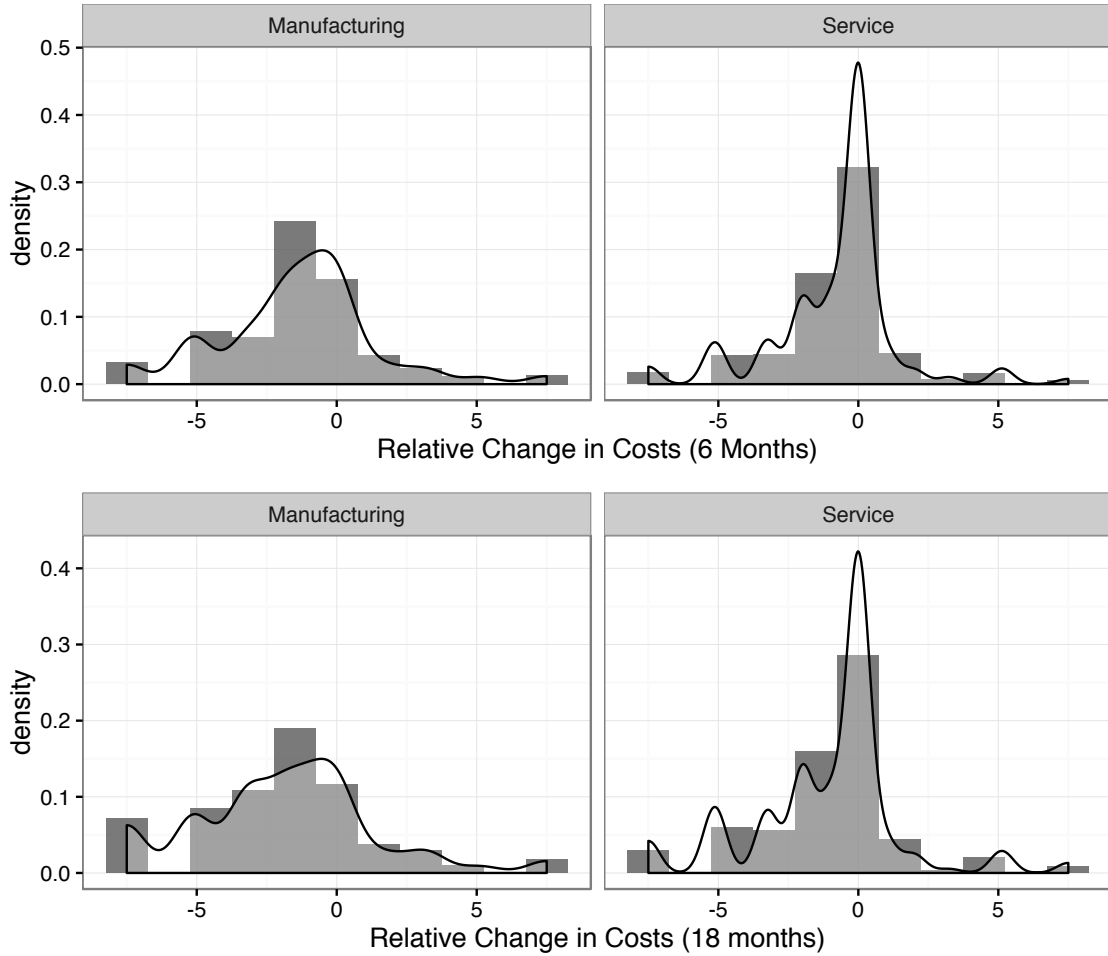
The histograms show the expected change in turnover for a 6-month ahead horizon as well as a 18-month ahead horizon. The bars show the relative frequency of firms' expectations ranging from $\leq -7.5\%$ to $\geq 7.5\%$. The upper row depicts expected changes for a 6-month horizon while the bottom row show the expected changes for an 18-months horizon. The shaded polygon represents the smoothed kernel density.

Formally:

$$\delta_{i,s} = \beta \mathbf{x}_i + \gamma \mathbf{z}_i + \psi \mathbf{d}_s + \xi_{i,s}$$

where $i = 1, \dots, I$ and $s = 6$ months, 18 months. $\delta_{i,s}$ is the survey-based impulse response in either turnovers, costs or profits of firm i at horizon s to the exchange rate

Figure 4.6: Cost distributions



The histograms show the expected change in costs for a 6-month ahead horizon as well as a 18-month ahead horizon. The bars show the relative frequency of firms' expectations ranging from $\leq -7.5\%$ to $\geq 7.5\%$. The upper row depicts expected changes for a 6-month horizon while the bottom row show the expected changes for an 18- months horizon. The shaded polygon represents the smoothed kernel density.

shock. \mathbf{x}_i is a row vector of firm-specific explanatory variables and \mathbf{z}_i is a row vector of J industry dummy variables where the j -th = 1st, ..., J -th dummy variable takes value 1 if firm i is in industry j and zero otherwise. \mathbf{d}_s represents a time dummy equal to 1 if $s = 18$ months and zero otherwise. \mathbf{x}_i includes the following variables: firm i 's size measured by its number of employees, firm i 's (non-shock scenario)

export shares to the Euro area and to the rest of the world, firm i 's (non-shock scenario) import share from the Euro area, firm i 's trade surplus with the Euro area and the rest of the world, firm i 's oil dependency defined as the share of expenditures on oil relative to total costs, firm i 's market power measured by its profit margin, (i.e. (total sales – total costs)/total sales, averaged over 2010–2012) and firm i 's unconditional Swiss franc/Euro exchange rate forecast for 2013. β is a row vector of coefficients attached to \mathbf{x}_i and γ can be seen as a vector of industry-specific intercepts or industry-specific fixed effects that control for unobserved heterogeneity between industries. ψ is the 18 months time fixed effect. A summary of all variables can be found in the appendix at the end of the thesis. $\xi_{i,s}$ is the error term. The regression coefficients are estimated by ordinary least squares.

Table 4.2 shows the estimation results of the aforementioned regression set-up. The first column shows that a higher Euro area export share is associated with a stronger reduction in expected turnover in response to the exchange rate shock. This effect is statistically significant at the 1% level. Also, the time fixed effect is negative and statistically different from zero at the 1% significance level. This statistical finding implies that the reduction in expected turnover, in response to the exchange rate shock, is stronger at the 18-months horizon than at the 6-months horizon. Conversely, higher market power appears to dampen the negative effect of an exchange rate shock. Though this result confirms the economic intuition, we restrain from putting too much weight on this particular finding because it is only significant at the 10% level. Firm size, proxied by the number of employees, does not have an effect on firm-level impulse responses. In addition, we find non-linear effects for firms' unconditional exchange rate expectations for 2013. In other words, the total marginal effect of the unconditional exchange rate expectation for 2013 depends on its level. Setting the derivative to zero yields $fxrate_exp = -\frac{\beta_{fx}}{2*\beta_{fxsq}} = 116$. This implies that firms which unconditionally anticipated the exchange rate to be at 1.16 Swiss franc per Euro do not report additional changes in expected turnover. Firms which anticipated the Swiss franc to be weaker than 1.16 tend to report an additional negative contribution to turnover impulse responses. However, the results for the turnover reaction are only statistically significant at the 10% level.

Table 4.2: Changes in turnover, costs, and profits given exchange rate shock

| | Turnover | Costs | Profit(6 months) |
|-------------------------|---------------------|---------------------|---------------------|
| Number of employees | 0.0001 (0.0001) | -0.0001 (0.0002) | 0.0003 (0.0003) |
| Market power | 0.6869 (0.4102)* | -0.1740 (0.3981) | 0.1932 (0.6055) |
| Year dummy | -0.6586 (0.1592)*** | -0.4156 (0.1536)*** | |
| Fxrate expectations | 0.5591 (0.2979)* | 0.5577 (0.2810)** | 0.8618 (0.4131)** |
| Fxrate expectations sq. | -0.0024 (0.0013)* | -0.0024 (0.0012)** | -0.0038 (0.0018)** |
| Export share euro | -0.0557 (0.0039)*** | | |
| Export share world | -0.0022 (0.0042) | | |
| Import share euro | | -0.0135 (0.0037)*** | |
| Import share world | | 0.0124 (0.0060)** | |
| Trade surplus euro | | | -0.0601 (0.0056)*** |
| Trade surplus world | | | -0.0103 (0.0062)* |
| R ² | 0.2879 | 0.0877 | 0.3772 |
| Adj. R ² | 0.2748 | 0.0695 | 0.3482 |
| Num. obs. | 1159 | 1075 | 450 |
| RMSE | 2.7094 | 2.5174 | 2.4022 |

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Models control for sector specific effects using sector dummies.

The second column of Table 4.2 shows the regression results for the cost impulse responses. Similar to firms' turnover reaction, the cost reaction is stronger for the 18-months horizon than for the 6-months horizon. Also, the import share for Euro area imports has a highly significant negative effect on firms' cost responses. A 10% higher import share leads to an additional total cost reduction of 0.1 percentage points. This effect is statistically significant at the 1% level. In contrast, a higher share of imports from the rest of the world dampens the cost reduction effect of the exchange rate shock as firms likely assumed other Swiss franc currency pairs were not affected by the scenario. We find that a 10% higher import share from the rest of the world dampens the cost response by 0.1 percentage points. Neither market power nor firm size have a significant effect on the cost impulse responses. Again, a closer look at the role of firms' unconditional exchange rate expectations for 2013 yields interesting results: Setting the derivative to zero results in $fxrate_exp = -\frac{\beta_{fx}}{2*\beta_{fxsq}} = 114.8$. Firms that expected an unconditional exchange rate of 1.148 Swiss franc per Euro for 2013 do not report additional changes in expected costs. Firms that anticipated

the Swiss franc to be weaker than 1.148 tend to report a more pronounced cost reduction. Both coefficients used to model the effect of unconditional exchange rate expectations are statistically significant at the 5% level for the cost model.

Finally, the third column contains estimation results for the profit impulse responses – i.e, the difference in profits between the baseline and our exchange rate scenario. The export and import shares are highly correlated, hence, using both variables in the regression would result in multicollinearity issues. In order to prevent these issues, we use trade surpluses with trading partners in the Euro area as well with partners from the rest of the world.⁹ Higher trade surpluses with partners from the Euro area are linked to additional profit reductions in the scenario. This effect is highly significant and more substantial than the effect for trading partners from the rest of the world. We also find a weaker negative effect of trade surpluses with partners from the rest of the world, which are only significant at the 10% percent level though. Market power and firm size do not have a statistically significant effect on profit impulse responses. We find a non-linear effect of unconditional exchange expectations in the profit model which is similar to the effects we found for turnover and costs. Firms that expected the average exchange rate for 2013 to be at 1.127 Swiss francs per Euro ($fxrate_exp = -\frac{\beta_{fx}}{2*\beta_{fxsq}} = 112.7$) tend to not report an additional effect on profits, while firms which expected the Swiss franc to be weaker than 1.127 tend to report a more pronounced profit reduction given the scenario. Again, both coefficients are statistically significant at the 5% level.¹⁰

4.4 Conclusion

In this paper, we have presented an application of survey-based impulse responses to analyze the effects of an exchange rate shock on Swiss firms' turnovers, total costs

⁹For a detailed explanation on the computation of trade surpluses in this paper, see appendix to this chapter at the end of the thesis.

¹⁰Note that the sample for the profit model is substantially smaller. This is because profit is a composite variable which requires turnover and costs not be missing, which reduces the available observations. However, more importantly, part of the data to compute trade surpluses was only available for 2012. Hence, we could only estimate the model for the 6 months horizon.

and profits. Survey-based impulse response analysis provides a convenient way to identify macroeconomic shocks without the need to impose parametric restrictions. Survey-based impulse responses can create structural microeconomic data allowing to shed light on the heterogeneity of economic agents and allow to easily capture any kind of non-linearities. Survey-based impulse response analysis is “on the spot”, i.e., it determines the effects of shocks at the time the survey was conducted rather than through historical time series. This feature makes the approach especially valuable in times of a structural break.

In July 2012, we applied a macroeconomic treatment scenario to a representative sample of nearly 900 Swiss firms. We asked firm representatives to evaluate the effects of a hypothetical change of the exchange rate floor from 1.20 to 1.10 Swiss francs per Euro – and a subsequent appreciation of the Swiss franc by the same magnitude – on expected firm-specific turnovers, total costs and profits 6 months and 18 months ahead. Our findings suggest an incomplete exchange rate pass-through to firms’ costs. Moreover, exchange rate shocks seem to be absorbed by firms given that their turnovers decrease more than their costs. There is also evidence that impulse responses are diverse across firms and industries. The manufacturing sector reacts more strongly than the service sector or the construction sector in the size of expected turnover reductions, as well as the size of expected total cost reductions. The industries being affected the most, within manufacturing, are machine producers & car suppliers as well as metal producers. At the same time, especially firms of these industries benefit the most from costs reductions in consequence of lower import prices, thereby dampening the potentially negative short-term effect on the economy caused by the exchange rate appreciation shock. Further, it turns out that the exchange rate shock leads to a substantial reduction in firms’ expected profits. While profit reductions are stronger for the manufacturing sector aggregate than for the service sector aggregate, we again find a high variation at the (sub-sector) industry level and the firm level.

Firm-level panel regression analysis allows us to control for unobserved industry effects. We find that a higher Euro area export share is associated with a stronger reduction in expected turnovers in response to the exchange rate shock. Moreover,

a higher import share from the Euro area is associated with a stronger reduction in expected total costs in response to the exchange rate shock. A higher trade surplus with the Euro area and the rest of the world leads to a more pronounced negative impulse response to the exchange rate shock.

Further, the degree to which firms would be caught off guard, measured by their unconditional exchange rate expectations for 2013, significantly relates to firms impulse responses. Interestingly, the additional effect of unconditional exchange rate expectations is zero at points (1.16, 1.148, 1.127) which are reasonably close to the scenario's 1.10 Swiss franc per Euro. Also, all three of these points are clearly below the 1.20 exchange rate floor that was in place when the survey was conducted. Firms that already expected a situation similar to the scenario tended not to report an additional effect of their own exchange rate expectations. Firms that anticipated a very strong Swiss franc in their unconditional expectations clearly dampened the negativity of their impulse responses. All three models show that firms reacted to the scenario and were able and willing to form conditional expectations in general.

Chapter 5

Ask Your Doctor or Pharmacist! On the Impact of Physician Dispensing on Utilization of Pharmacies*

*This chapter is based on KOF Working Paper 387.

5.1 Introduction

In most advanced economies patients get their medicines either directly from their doctors (self-dispensation) or via a pharmacy. The superiority of either distribution system continues to be a lively political and economic debate: the general idea of splitting prescription and sale of drugs is to avoid the incentive to over-prescribe or to prescribe expensive drugs instead of equally effective cheaper medication (e.g. generics). However, while the separation of prescription and sale of drugs is reasonable in most regions, factors such as topography and population density can justify self-dispensation. Although there is evidence that self-dispensation increases general drug expenditures (Kaiser and Schmid, 2014), better medical and pharmaceutical coverage can outweigh these additional costs.¹ Yet, there are other possible effects of physician dispensing that cannot be seen by solely looking at total health care expenditures. The quality of pharmaceutical and medical services may differ when the service is provided by a single individual as opposed to two distinguished specialists.

In this paper we intend to supplement the discussion on physician dispensing, which so far has mostly focused on total health care expenditures. Our view instead concentrates on the impact of physician dispensing on the orientation of the pharmaceutical and medical professions themselves. We study pharmacy-level data from Switzerland in order to quantify the impact of physician dispensing on the demand of pharmaceutical services from pharmacies, while controlling for pharmacies' socio-economic and spatial environment.

By relating these effects to shifts in pharmacies' portfolios as well as to the overall size of the pharmaceutical market, we intend to show that the substitution effects are substantial and that physician dispensing may influence the health care system beyond its impact on total health care expenditures. For example, a shift in the orientation of physicians and pharmacists would influence accessibility of drugs substantially. Physicians generally only provide their own patients with drugs and have different opening hours and infrastructure than pharmacies.

¹Section 5.2.1 lists further studies of the impact of self-dispensation on health expenditures.

Switzerland is an ideal case for our study, because the country's federalist legislation allows us to investigate different dispensation regimes within the same country. Also, the country's long tradition of self-dispensation exposes long-term structural differences in the number of dispensing physicians and pharmacies between regions with diametric legislation. Using pharmacy-level data, we find substantial effects on pharmacy usage and pharmacies' portfolios beyond these established structural differences. This is important for policy makers inasmuch as if physician dispensing is implemented, beyond the idea of filling coverage voids, it may lead to a two-tier system. In turn a two-tier system may absorb shocks, such as changes in drug price regulations, differently across tiers. Hence, this paper encourages decision makers to monitor consequences of comprehensive physician dispensing beyond its effects on total health care expenditures. Careful monitoring of all dispensation channels is essential to fully assess the effects of shocks such as changes in drug price regulation.

The remainder of the paper is structured as follows. The second section discusses the legal situation in Switzerland and provides an overview of the pharmaceutical market as well as an overview of the literature on self-dispensation. The third section introduces our dataset. The fourth section covers our methodological approach including an assessment of coverage as well as our regression setup. Section 5 presents our estimation results and, lastly, section 6 concludes and summarizes our findings.

5.2 Physician dispensing

This section consists of three subparts: The first part gives a general overview of different forms of drug dispensation that can be found in the literature. The second part focuses on self-dispensation and elaborates on the specific legal situation in Switzerland. The third part gives an overview of the Swiss pharmaceutical market.

5.2.1 General overview

The terms and conditions of drug dispensation have historically been subject to a heated debate between two professions. Trap (1997) dates the dispute between pharmacists and doctors back to France in the 13th century. A few years earlier

the German emperor Fredrick II initiated the separation of the two professions and created the basis of the current system in most European countries. Recently, several scholars studied the effects of self-dispensing physicians with particular focus on the effect of different regimes on health expenditures. Although current research does not vote unambiguously in favor of either system, the aforementioned strand of literature remains an important motivation for our research as nearly all studies find differences between the two channels. It is important to monitor which drug dispensation channel patients choose, because patients and their doctors usually do not bear the full costs of their decision. Thus neither of them has an incentive to choose the more cost-effective channel if the other channels are more convenient or promised additional income.

In most developed countries, doctors are not allowed to sell drugs directly to their patients (Filippini et al., 2013). Several studies focus on Switzerland because the Swiss legislation allows the study of both self-dispensation and non-self-dispensation regimes within the same country. Compared to pharmacist dispensation, Kaiser and Schmid (2014) find that physicians in Switzerland produce higher drug expenditures than pharmacists in the order of 30% per patient. Likewise, Beck et al. (2004) find higher drug expenditures in Swiss cantons that allow self-dispensation, correcting for socio-economic variables. Rischatsch and Trottmann (2009) show that self-dispensing doctors in Switzerland have a higher probability of prescribing the drug with the higher margin compared to non-dispensing doctors. Busato et al. (2010) examine whether treatment costs differ across medical disciplines and regimes for the years 2003-2007. Depending on the professional discipline their results vary: they find significant arguments for and against lower costs of either regime. For the most expensive treatment (by non-invasive specialists) they find significantly lower costs in the prescription only case.² Reich et al. (2012) show that an increase in the density of dispensing doctors leads to an increase in per capita health care expenditures.

In a recent study commissioned by the Swiss Federal Office of Health, Trottmann et al. (2015) examine whether patients in cantons that allow self-dispensation have

²They also take regimes with non-explicit legislation into account.

the same level of drug consumption as patients in cantons without self-dispensation.³ Comparing two similar cantons and correcting for socio-economic variables they find lower *ceteris paribus* drug expenditures under self-dispensation and a higher likelihood of prescribing generics. At the same time, self-dispensation leads to higher expenditures for medical services (consultations). Overall, Trottmann et al. (2015) do not find differences in the level of consumption of health services between both dispensing regimes. Their analysis refers to services granted under the compulsory health insurance.

Most international studies find self-dispensation to increase health care expenditures. For Japan, Iizuka (2007) shows that the additional markup of selling drugs affects doctors prescription choices. Physicians tend to over-prescribe, and, as a second effect might not choose the optimal medicine from a patient's perspective. In line with the findings of Iizuka (2007), Liu et al. (2009) add that physicians' prescription decisions depend on the profit margin between the reimbursement and the acquisition price. This moral hazard of physicians has also been described by Lundin (2000) who nevertheless shows that, despite physicians' knowledge about price differences, physicians do not strictly act based on this information.

Based on the reduction of drug expenditures in Taiwan after self-dispensation was banned, Chou et al. (2003) claim that self-dispensation increases expenditures for medicine on a per visit basis. In a systematic literature review Emery et al. (2009) examine 21 papers on the comparison of self-dispensing and non-dispensing doctors' practices. The examined studies cover the U.S. (6 papers) and the UK (5), Zimbabwe (5), South Korea (2), Australia (1), South Africa (1), and Taiwan (1). Emery et al. (2009) conclude that self-dispensing physicians tend to prescribe more pharmaceuticals, produce higher pharmaceutical costs and are less likely to prescribe generics than non-dispensing doctors.

Furthermore, studies such as Greer and Jacobson (2010) or Uhlmann (2013) which focus on the aspect of federalism in health care politics, are relevant for our

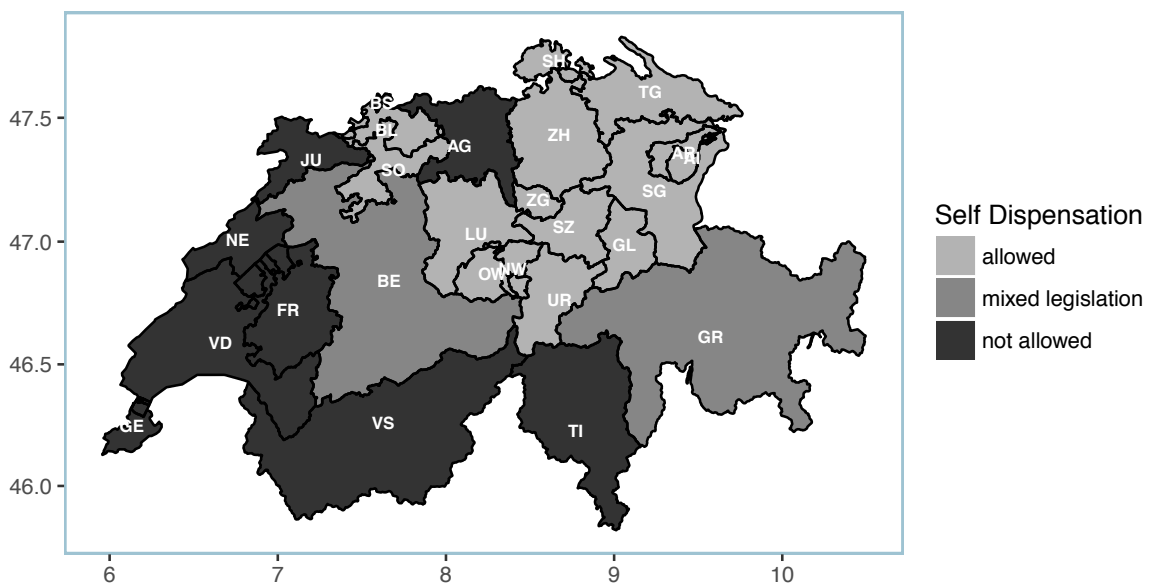
³The study was finished in 2014, but was publicly available only from 2015 onwards.

analysis. Though all of Switzerland is technically a federalist country, the cultural influence from federalist Germany and centralistic France on the German- and French-speaking parts of Switzerland, respectively, can be seen on many occasions. In our case, the French-speaking part has a unitary ruling while the German-speaking cantons make use of the federalist structure and allow for differences in their legislation.

5.2.2 State of legislation in Switzerland

Switzerland is divided into 26 states (cantons) and has a long tradition of organizing many aspects of legislation at the state level. This is also the case for the legal parameters of drug dispensation.⁴ Hence, different regimes can be found in Switzerland. Figure 5.1 gives an overview of the current legislation across Swiss cantons.

Figure 5.1: Legislation by Canton



⁴The Swiss law on health care insurance states in article 37, letter 3 (Version of 1 July 2013): "Die Kantone bestimmen, unter welchen Voraussetzungen Ärzte und Ärztinnen mit einer kantonalen Bewilligung zur Führung einer Apotheke den zugelassenen Apothekern und Apothekerinnen gleichgestellt sind. Sie berücksichtigen dabei insbesondere die Zugangsmöglichkeiten der Patienten und Patientinnen zu einer Apotheke."

The dark areas show cantons that do not allow physicians to dispense drugs. The light gray areas illustrate cantons that allow doctors to prescribe and hand out drugs in general. While the medium shading highlights the two cantons that have mixed legislations: Berne and Grisons. Mixed legislation refers to a non-explicit legislation and means that the actual ruling may differ on the municipality level.⁵ Because such situations can be very specific and hard to compare our empirical analysis will only consider cantons that have an explicit legislation. In general, all of the French-speaking western cantons and Ticino form a group of cantons that prohibit self-dispensation. The German speaking part is less unified, but is generally more open to physician dispensing.

Note that, until 2012, self-dispensation was prohibited in the city centers of Schaffhausen (SH), Winterthur (ZH) and Zurich (ZH). Though physicians were allowed to dispense in the respective cantons. In our paper, we account for these exceptions with the help of zip codes and assign the relevant legislation to all affected pharmacies. Legislation is technically time-invariant in our study. Though several referenda took place between 2008 and 2012, no regime changes that would have allowed reasonable intertemporal comparisons within the same canton, took place during the observed period. In fact, the aforementioned exceptions for the city centers of Schaffhausen, Winterthur and Zurich were abolished in May 2012. Yet, the changes became effective too late in the year 2012 to consider 2012 under self-dispensation for the affected pharmacies. For the year 2012, we rather consider Schaffhausen, Winterthur and Zurich as we did for 2011, not only because the change came in the middle of the year, but also because the exact implementation date was difficult to anticipate due to the lengthy legal process. Given the lengthy adaption phase, we continue to consider these pharmacies to be located in a region that prohibits physician dispensing. Yet, intertemporal comparisons for 2013 and beyond would be an interesting subject for further research, but are beyond the scope of this paper. However, the fact that larger cities such as Zurich and Winterthur introduce

⁵Most cantons do have single exceptions, i.e., physicians with exceptional licenses to dispense drugs. This is the case in almost every canton and is not referred to as mixed legislation. Mixed legislation refers to a situation in which legislation varies across the canton. Single exceptions are ignored as they do not influence the aggregate effects.

self-dispensation is noteworthy as there are obviously other reasons to introduce self-dispensation than improving pharmaceutical coverage in remote villages, a reason that has little basis in populated urban centers.⁶

5.2.3 Pharmaceuticals market in Switzerland

Table 5.1: Pharmaceuticals Market in Switzerland, source: interpharma.ch

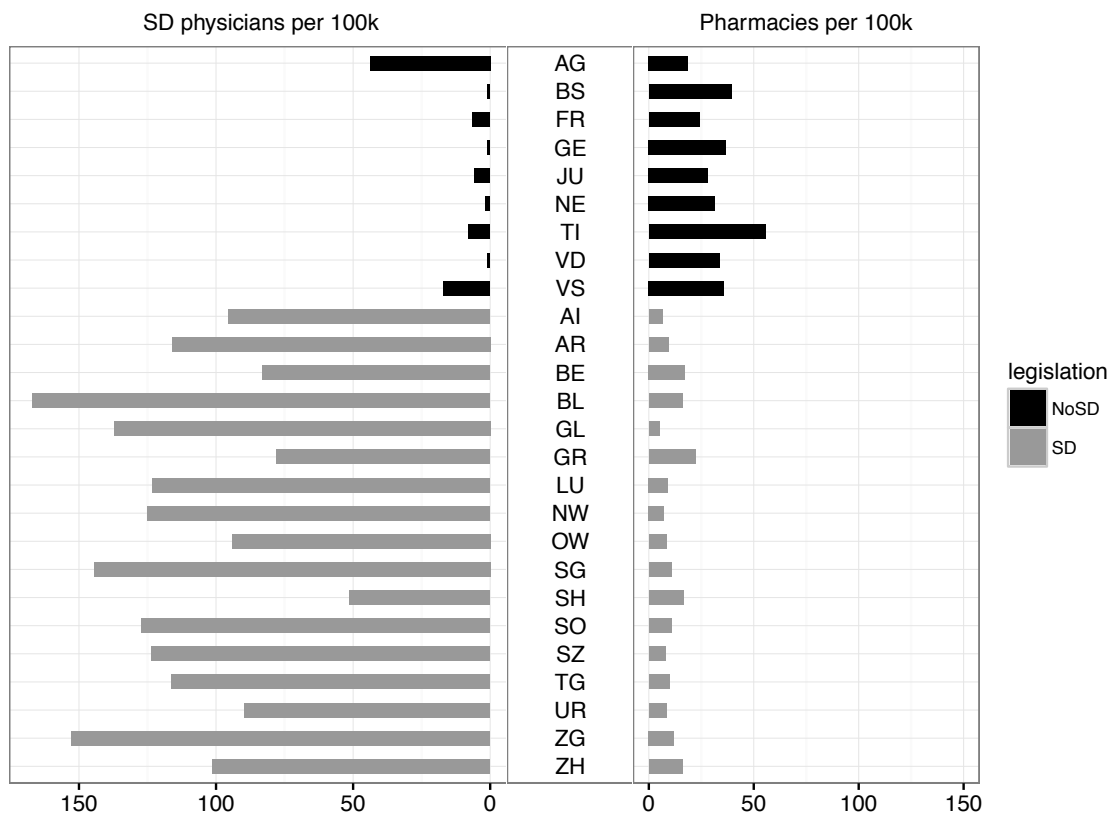
| Channel | Volume (CHF) | Market share | Packages sold | Cost per package |
|---------------|--------------|--------------|---------------|------------------|
| Pharmacies | 2649.50 | 0.52 | 116.40 | 22.76 |
| SD physicians | 1233.70 | 0.24 | 38.50 | 32.04 |
| Hospitals | 1122.30 | 0.22 | 43.80 | 25.62 |
| Drug Stores | 77.00 | 0.02 | 8.90 | 8.65 |
| Total | 5082.60 | 1.00 | 207.60 | 24.48 |

Table 5.1 shows how the volume of the Swiss pharmaceutical market is shared among different distributors. At 2.65 billion CHF, pharmacies account for 52% of the total market volume (5.08 billion CHF). In 2012, dispensing physicians represented more than 24% of the market volume and sold pharmaceuticals worth more than 1.2 billion CHF. Further, the cost per package sold by dispensing physicians is substantially higher (32.04 CHF per package) than the cost per package sold by pharmacies (22.76 CHF per package). Although the above data are reported in terms of manufacturing prices and thus not directly comparable to the pharmacy-level revenues analyzed in our survey, the data illustrate the size and market shares of the pharmaceutical market and its distributors in Switzerland.

Based on the above numbers, in 2012, an average pharmacy sold drugs worth 1.56 million CHF (2.65 billion CHF / 1740 pharm.) and an average dispensing physician sold drugs worth 0.21 million CHF. The self-dispensing physicians' total market share of 1.2 billion CHF corresponds to the revenue of more than 770 average pharmacies.

⁶A referendum to change the current legislation and allow physicians to dispense drugs was turned down in the canton of Aargau in September 2013. Although the referendum was turned down, the fact that a referendum was initiated may indicate that there is a lobby component to the dispensation debate. For a non-self-dispensing canton, the canton of Aargau has an extraordinarily high number of dispensing physicians who operate as single exceptions. See Figure 5.2.

Figure 5.2: Dispensing physicians and pharmacies by Canton
 source: interpharma.ch



The substantial market share of physicians has led to the establishment of structural differences in the landscape of drug dispensers across physician dispensing and non-physician-dispensing cantons. Figure 5.2 gives a canton-by-canton overview of the number of pharmacies and dispensing physicians per 100 thousand inhabitants. Self-dispensing cantons clearly have fewer pharmacies per capita: on average, cantons that prohibit self-dispensation have 32.9 pharmacies per 100 thousand inhabitants while self-dispensing cantons only have 12.9 pharmacies per 100 thousand inhabitants. In other words, given that self-dispensing cantons account for a population of 3.64 million people and non-self-dispensing cantons for 3.21 million people, one would need approximately 730 additional pharmacies in self-dispensing cantons to level the coverage gap between both types of cantons.⁷ At 770 pharmacies, physicians total revenue slightly exceeds the benchmark of 730 pharmacies. This difference may indicate that total expenses are slightly higher under a self-dispensation regime. However, the gap between physicians market share and the benchmark is moderate and could also be partly the result of topographic differences.

Nevertheless, we can immediately see substantial structural differences in the number of pharmacies across cantons and regimes, respectively. Yet, this does not elucidate whether structural differences are the only consequence of physician dispensing on pharmacies. Changes in the utilization of pharmacies could also alter pharmacies' portfolios and thus influence the accessibility of drugs as well as the quality of pharmaceutical services.

5.3 Data

The analysis in this paper is centered around a Swiss pharmacy-level dataset, but uses data from multiple sources to enrich the original dataset. In this section we first describe all variables used in our analysis. Second, we provide information on our sample of pharmacies and the distribution of focal variables.

⁷Note that we do not consider the cantons of Grisons and Berne because of their mixed legislation. Berne and Grisons had 1.19 million inhabitants in 2012.

5.3.1 Variables

Table 5.2 describes all variables used in our analysis. The analysis is centered around a panel survey study called *RoKA*. *RoKA* is conducted on a yearly basis by the KOF Swiss Economic Institute on behalf of pharmaSuisse.⁸ Table 5.2 denotes variables stemming from this survey as *RoKA*. Revenue, as a proxy for a pharmacy’s utilization rate and the ratio of non-drug related revenue to total revenue, as an indicator for a pharmacy’s orientation, are particularly important to our study.⁹

We enrich pharmacy-level data by geo-location information obtained from Google’s Geocoding and Places Application Programming Interfaces (API). Further, we match municipality-level information provided by the Swiss Federal Statistical Office (SFSO) to our dataset. Besides standard demographic municipality information, we match a typology of municipalities to our dataset.¹⁰ This additional typology information allows to classify a pharmacy’s environment as a city center, agglomeration, mixed area or truly rural area. Although the SFSO offers more finely grained typologies with up to 22 categories, we favor the four groups mentioned above. At this level of granularity, matching municipality types to the dataset works reasonably well with the help of ZIP codes contained in the pharmacies’ addresses. Table 5.2 gives an overview of all the variables we use in this paper.

⁸pharmaSuisse is the association of Swiss pharmacies. pharmaSuisse commissions the cost focused *RoKA* (Rollende Kostenstudie der Apotheken) study on a yearly basis.

⁹Quantitative surveys that involve larger and / or sensitive numbers often suffer from quality issues as participants are reluctant to state precise numbers or fail to enter numbers with more than a handful of digits thoroughly. Using variables computed from multiple original variables increases the likelihood of item non-response and errors in these computed variables. Total revenue is a well-received question in the *RoKA* survey and suffers much less from occasional drop-out errors than more finely grained accounting-related questions. Because only one single question is involved, the data quality of the revenue variables is much higher than for, e.g., profit, which would have to be computed from multiple accounting-related questions.

¹⁰Municipalities that cannot be matched unambiguously by ZIP codes were adjusted according to a search on Google Maps.

Table 5.2: Variable description

| Variable | Source | Description | Unit |
|--------------------------|--------|--|--|
| Legislation | RoKA | physician dispensing indicator | categorical: (SD, NoSD) |
| Revenue | RoKA | total yearly revenue of a pharmacy | CHF |
| High VAT ratio | RoKA | non-drug related revenue : total revenue | share |
| Lng | Google | longitude | degree |
| Lat | Google | latitude | degree |
| Nearest_3 | KOF | avg. distance to 3 closest pharmacy | kilometers |
| Nearest_5 | KOF | avg. distance to 5 closest pharmacy | kilometers |
| Total area | RoKA | size of a pharmacy | square meters |
| Chain | RoKA | is pharmacy affiliated to a chain? | logical: true, false |
| Wkly. business hours | RoKA | business hours per week | numerical |
| Region | SFSO | type of municipality | categorical: city center, agglomerations, mixed, rural |
| Share of people aged 65+ | SFSO | people over 65 | ratio |
| Settlement area | SFSO | share of settlement area | ratio |
| Median tax inc | SFSO | median taxable income / 1000 | CHF |
| Legal form | RoKA | legal form of pharmacy | categorical |
| Pharmacist status | RoKA | is the pharmacist owner or co-owner of the pharmacy? | categorical |
| Year dummy | RoKA | year indicator | categorical |
| Wage included | RoKA | does RoKA cost report include pharmacist wage? | logical |

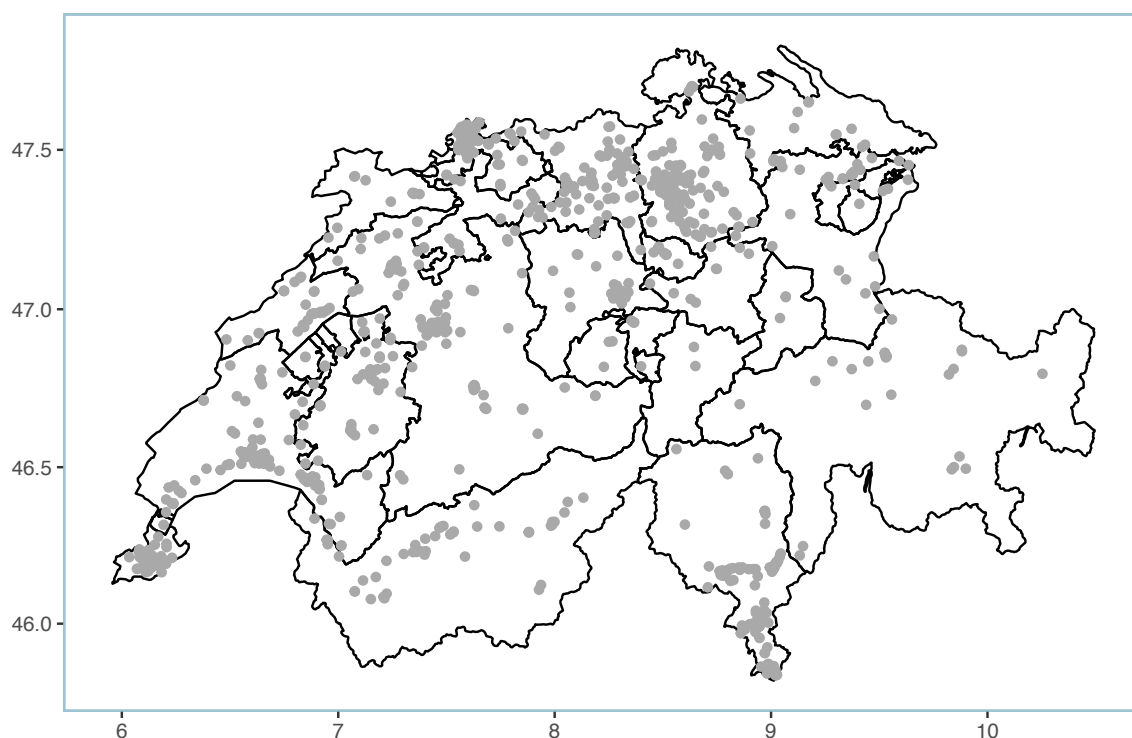
5.3.2 Sample

The *RoKA* study is conducted among all pharmaSuisse-pharmacies. And participation in the *RoKA* study is mandatory for members of pharmaSuisse. Currently about 77% of all Swiss pharmacies (1744) are affiliated with pharmaSuisse (pharmaSuisse, 2014). Given a response rate of more than 59%, 1038 pharmacies have regularly taking part in the survey between 2008 and 2012.¹¹ Figure 5.3 depicts the locations of all participating pharmacies. We can easily spot densely populated city

¹¹Regularly refers to taking part at least four out of five times.

centers and agglomerations of the country's largest cities, but we can also identify elements of Swiss topography, such as the Rhone Valley (south between the 7th and 8th degree of longitude) or Lake Geneva. Note that the cantons of Berne and Grisons are excluded from our analysis because these cantons have mixed jurisdiction and we only aim at comparing explicit dispensation regimes. We remove observations with unreasonable values as well as extreme outliers from the dataset. Specifically, we leave out all observations with more than 15 million CHF in revenue, more than 168 or less than 10 business hours per week. These corrections only remove about 20 observations per year, and, after typology matching and the exclusion of Berne and Grisons, leave us with 622 pharmacies, 473 of which are located in cantons that prohibit self-dispensation. Consequently, 149 pharmacies in our sample operate in cantons that allow physician dispensing.

Figure 5.3: Locations of RoKA study participants



It is important to keep in mind that self-dispensation has a long tradition in Switzerland and that we are comparing two regimes which have been established

long before the *RoKA* survey was conducted. Given that legislation is time-invariant in our dataset, we do not expect any substantial intertemporal dynamics. Figure 5.4 shows the number of pharmacies per 1000 inhabitants grouped by regional typology. As expected, we cannot speak of time trends – particularly if we disregard mixed and rural areas due to their small sample size. Yet, we do observe substantial level differences: regions that allow physician dispensing have substantially fewer pharmacies per 1000 inhabitants across all region types.

Figure 5.4: Pharmacies per 1000 citizen by regional typology (2012)

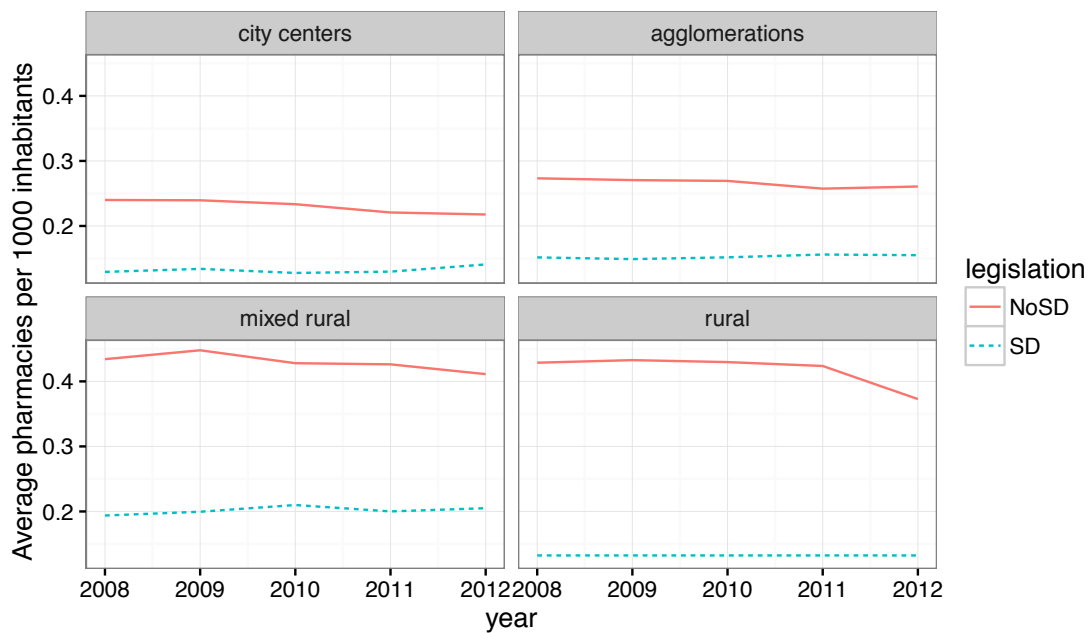
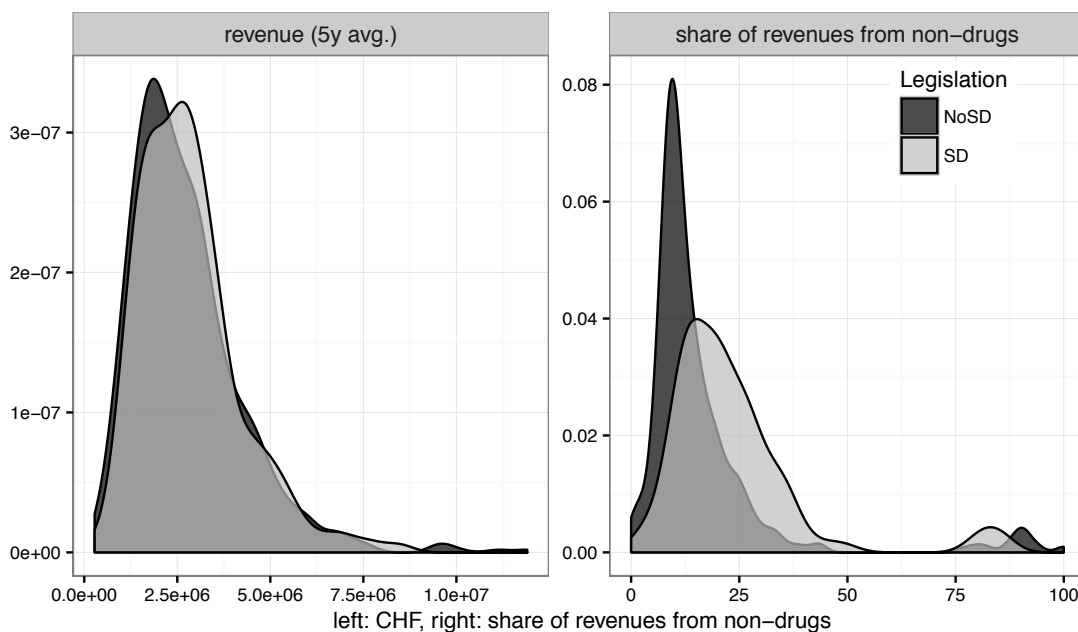


Figure 5.5 shows the density of total revenue and ratio of revenues with non-drugs respectively. We use both variables as dependent variables in our regression setup (section 5.4). While we cannot infer a meaningful difference in total revenues between the two regimes at first glance, the distribution of non-drug related revenues seems to differ substantially.

Figure 5.5: Density of dependent variables



5.4 Empirical strategy

Recall that we are interested in two related research questions: First, we want to explore whether patients in areas with physician dispensing substitute pharmacies with self-dispensing physicians. Second, we intend to measure whether pharmacies' orientation differs across the legislation settings. Further, note that even though we have a panel dataset at hand, legislation is a time-invariant variable within the observed timespan. Therefore, with the available data, our strategy is to pool data and control for basically time-invariant pharmacy-level variables as well as for municipality-level information. Again, we split this section into two subsections: The first subsection discusses the models used to address both questions. The second subsection briefly discusses how we measure coverage in a particular area.

5.4.1 Utilization of pharmacies

We use a binary legislation variable to predict pharmacies' revenues while controlling for a set of pharmacy- and municipality-level variables to account for the environment and circumstances in which a pharmacy operates. Because, we only consider cantons with explicit legislation with respect to self-dispensation, excluding Berne and Grisons from our analysis, we can use a binary indicator in our model which equals one when a canton allows physicians to dispense drugs and zero otherwise.¹² More formally we estimate the following pooled model using OLS:

$$Y = L'\beta + P'\gamma + M'\mu + T'\tau + \epsilon$$

Where Y denotes pharmacies' log revenues and L is a binary indicator for legislation. P is a matrix of pharmacy level control variables such as *opening hours per week* or mean distance to the closest competing pharmacies. M denotes a set of municipality-level variables such as *median taxable income* or *share of cultivated area*. Finally, T denotes a set of time dummies to account for yearly idiosyncrasies, such as business cycle movement in the aftermath of the crisis or pronounced flu seasons.

To study the second question mentioned in the previous section, we exploit the fact that drugs are subject to lower VAT than other products. The most recent waves of the RoKA survey ask pharmacists to split their revenues based on the two groups of VAT. Thus we are able – at least for part of the data – to distinguish drug based revenue from other revenue. Instead of total revenues, we take the revenue subject to high VAT as a dependent variable in an otherwise analogous regression set-up.

5.4.2 Measuring coverage

We measure the coverage in a particular region by computing the distance between a pharmacy and all other pharmacies from the pharmacies' geo-locations. With the

¹²see also figure 5.1 in section 5.2.2.

respective longitude and latitude of two locations, distances can be computed using the Great Circle Distance. Simple trigonometric procedures that assume the earth to be spherical can already produce reasonable results. Considering an equatorial axis and a flattening factor makes computation of distances more complex, but accounts for an ellipsoid earth model. The pioneering approach of Vincenty (1975) suggests an iterative procedure, which became the basis of many of today’s state-of-the art methods. However, we use an approximation suggested by Meeus (1999), which produces very accurate results and does not rely on iteration.¹³ Approximation requires setting constants for the flattening factor f and equatorial radius r . The values of these constants ultimately depend on the selected ellipsoid model and choosing the World Geodetic System Standard WSG84 (NIMA, 2000) implies the following values:

$$r = 6378.137 \tag{5.1}$$

$$f = 1.0/298.257223563 \tag{5.2}$$

Following Meeus (1999), we can compute the distances $d_{i,j}$. We use these distances to measure a region’s pharmaceutical coverage as follows. Suppose,

$$D = \begin{pmatrix} 0 & d_{1,2} & \cdots & d_{1,n} \\ d_{2,1} & 0 & \cdots & d_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ d_{n,1} & d_{n,2} & \cdots & 0 \end{pmatrix}$$

is a matrix of distances, where the j -th column is a vector of distances from the j -th pharmacy to all pharmacies in the sample. Hence, the diagonal elements of the matrix D are naturally zero, containing the distance from a pharmacy to itself. In order to construct a metric for coverage, we use the distance matrix D to aggregate

¹³The Meeus distance is implemented in several R packages. We chose the implementation provided by *sp* (Bivand et al., 2013), (Pebesma and Bivand, 2005) because of its use of C++. Using C++ considerably speeds up computation of our distance matrix.

the distance of the c closest pharmacies to the j -th pharmacy. Suppose that d_j is a sorted sequence of the values of the j -th column of D :

$$d_j = \langle d_{j,1}, d_{j,2}, \dots, d_{j,n} \rangle, \quad \text{where } d_{j,1} \leq d_{j,2} \leq \dots \leq d_{j,n}$$

We can then easily compute the average distance to the c closest pharmacies to obtain a pharmacy level coverage metric C_j . Note that $k = 1$, is left out as the first element of d_j is always zero:

$$C_j = \frac{1}{n} \sum_{k=2}^c d_{k,j}$$

We can further generalize our coverage indicator by:

$$C_j = f(d_j, k), \quad \text{where } k = \langle 1, 2, 3, \dots, c \rangle.$$

Typically we use the median or the mean as a function to aggregate the truncated vector of distances. For example, if $c = 3$ was set, we would obtain the average or median distance of the j -th pharmacy to the three closest competitors. Airline distance is a sufficient measure of coverage because we focus on city centers and agglomerations in our analysis.¹⁴

5.5 Results

This section presents the estimations results of the models described in the previous section. We focus on two questions: first, we present the effects of self-dispensation on the utilization of pharmacies proxied by pharmacy revenues. Second, we show the effects of self-dispensation on pharmacies' portfolios. In addition, we inspect

¹⁴Google also offers a service to compute distances between two locations based on roads, but only allows for a limited number of requests per 24 hours. Hence, requesting an entire distance matrix would be time consuming. Also, due to our focus on city centers and agglomerations, we prefer a simple and transparent method over a third party black box computation. If one wants to put more emphasis on the importance of routes as opposed to airline distances, the approach of Luxen and Vetter (2011) in combination with *Open Streetmap* would be a promising starting point.

regional coverage visually to get an idea of how differing cantonal legislation may have established different long-term location patterns that cannot be represented within the scope of the survey.

5.5.1 Estimation results

Table 5.3 displays the estimation results of our pooled OLS models.¹⁵ The first column shows the regression coefficients for total revenues in the full sample. The second column presents estimation results for total revenues in a reduced sample that only considers pharmacies located in city centers and agglomerations. Finally, the third column presents the estimation results for revenues made with products, which are not subject to the reduced drug VAT. We exploit the distinction in VAT rates to distinguish drug revenue from non-drug related revenues. We find almost identical results using the full sample and a reduced sample which only considers more densely populated areas. We find a highly significant and substantial negative effect (15% less) of allowing physicians to dispense drugs on the utilization of pharmacies in both samples. We also find that non-drug related revenues are almost 40 percent higher in regions that allow physician dispensing.

The average distance variable measures regional coverage in order to control for competition by pharmacies in close proximity. We also add a quadratic term of this average distance to our model in order to model the non-linear influence. As expected the signs of both terms differ, indicating that not being strongly exposed to competing pharmacies increases utilization of a pharmacy in densely populated areas. Yet, a low density of pharmacies may lead to a decreased awareness of pharmaceutical products and services in a particular region.

We also control for the total size of a pharmacy, measured in square meters, to account for its size. We chose size instead of employees, because size is not volatile over the timespan of our survey. Larger pharmacies have significantly higher revenues: On average, a 10 square meter increase in area leads to a 1.5 percent

¹⁵The difference in the number of observation between the first two columns and the third column stems from the fact, that the dependent variable for the latter is only available for 2012 while total revenue is available for multiple periods.

Table 5.3: Utilization of Swiss Pharmacies

| | Revenue | Revenue (no rural) | High VAT Revenue |
|-------------------------------|---------------------|---------------------|--------------------|
| Intercept | 12.2203 (0.2261)*** | 12.2282 (0.2395)*** | 9.5159 (1.2598)*** |
| Physician dispensing allowed | -0.1564 (0.0203)*** | -0.1635 (0.0217)*** | 0.3704 (0.0873)*** |
| Average distance (3 closest) | 0.0497 (0.0071)*** | 0.0658 (0.0119)*** | 0.0494 (0.0353) |
| Average distance sq. | -0.0022 (0.0004)*** | -0.0040 (0.0011)*** | -0.0022 (0.0020) |
| Total area | 0.0015 (0.0001)*** | 0.0015 (0.0001)*** | 0.0009 (0.0004)* |
| Affiliated to chain | 0.0459 (0.0221)* | 0.0182 (0.0239) | -0.1472 (0.2022) |
| Wkly bus. hours | 0.0434 (0.0064)*** | 0.0405 (0.0064)*** | 0.0544 (0.0371) |
| Wkly bus. hours sq. | -0.0002 (0.0000)*** | -0.0002 (0.0000)*** | -0.0002 (0.0003) |
| Region: agglomeration | 0.0736 (0.0177)*** | 0.0617 (0.0190)** | 0.1570 (0.0955) |
| Region: misc. rural | -0.1163 (0.0257)*** | | -0.0765 (0.1406) |
| Region: rural | -0.0329 (0.0542) | | 0.1167 (0.1275) |
| Pop. share 65+ | -0.0040 (0.0025) | -0.0013 (0.0031) | -0.0026 (0.0131) |
| Share settl. area | 0.0008 (0.0004)* | 0.0008 (0.0004)* | -0.0000 (0.0022) |
| Med. tax. inc. (per 1000 CHF) | 0.0086 (0.0010)*** | 0.0086 (0.0014)*** | 0.0018 (0.0050) |
| R ² | 0.3377 | 0.3159 | 0.2114 |
| Adj. R ² | 0.3336 | 0.3112 | 0.1852 |
| Num. obs. | 3886 | 3198 | 622 |
| RMSE | 0.4416 | 0.4478 | 0.9474 |

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$. Models control for ownership status, legal form of the pharmacy, year dummy.

increase in revenues. We also control for affiliation of pharmacies to a chain of pharmacies. Chain affiliation has a positive effect on revenue in both samples, but it is only significant when excluding rural regions.

Our estimation results for weekly business hours are intuitive: Additional hours lead to higher revenue, though the gain is not linear. Very much like self-dispensation, business hours are subject to legal regulations and thus need to be controlled for in order to disentangle the effect of dispensation legislation.¹⁶

Being located in an agglomeration as opposed to a city center – other variables held constant – is associated with seven percent higher revenues. The coefficient of mixed rural areas and rural areas is negative but only significant for mixed rural areas. Although we might expect an effect of the share of citizens older than 65 years, we do not find a significant effect on the utilization of pharmacies. The negative coefficient is not necessarily intuitive, but might be caused by other structural aspects that are correlated with a high share of older people. The share of settlement area in a municipality has a significant positive effect on pharmacies revenues. Finally,

¹⁶We have also investigated the role of interaction terms between legislation and business hours, but did not find a significant interaction term.

we control for wealth in pharmacies environment by using median taxable income as a control variable. As expected, the revenue of pharmacies in wealthier regions is significantly higher, because many products in a modern pharmacy's portfolio are not subject to highly inelastic demand.

Allowing physicians to dispense drugs leads to a substantial decrease in the utilization of pharmacies when controlling for pharmacy- and municipality-level variables. Without controlling for the aforementioned variables we cannot find a pronounced effect of self-dispensation. Pharmacies face a reduced role as drug dispensers in self-dispensing cantons. However, the third column of Table 5.3 shows a *ceteris paribus*, 37 % increase in revenues from non-drugs for pharmacies in self-dispensing cantons. This may indicate that pharmacies, which are exposed to self-dispensation, are adapting their orientation and portfolios to their shifted role.

5.5.2 Regional coverage patterns

Although the results presented above quantify the impact of self-dispensation, their scope is limited to the current impact on pharmacies. Our regressions cannot display long term dynamics such as location decisions of pharmacies in the long run. Thus we intend to add a brief visual insight to the discussion. Visual inspection may help elucidate how Switzerland established two systems to provide patients with drugs and pharmaceutical services.

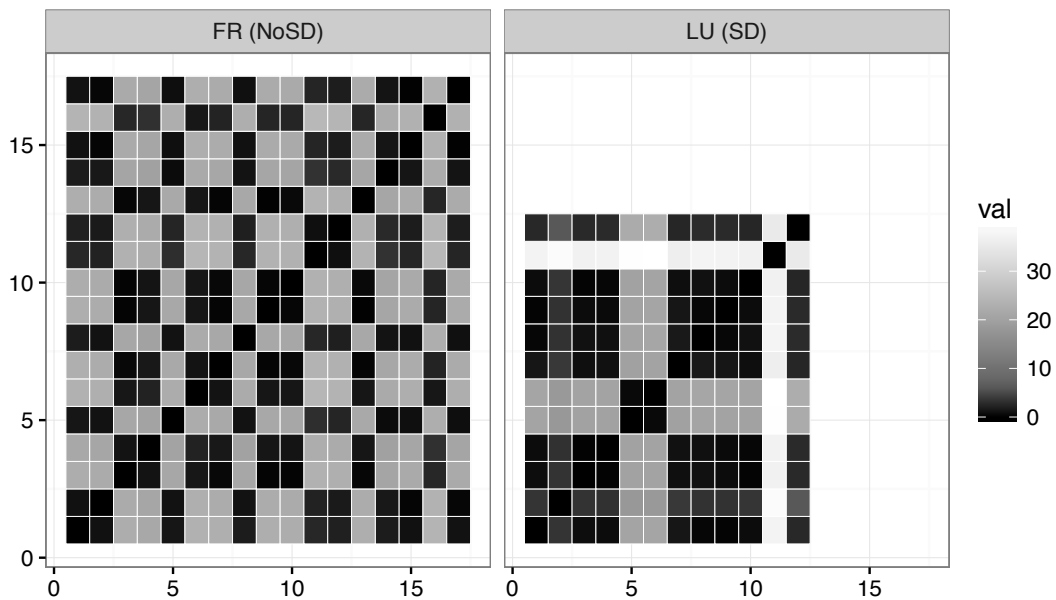
Figure 5.6 shows a comparison of two heatmaps indicating local concentration of pharmacies. The heatmap on the left shows distances between pharmacies in the city centers of the canton of Fribourg where self-dispensation is prohibited and right half shows the city centers of the canton of Lucerne where self-dispensation is allowed. The city centers of Lucerne and Fribourg can be considered topographically similar.¹⁷ In both halves, every square represents the distance between a pair of pharmacies. Hence, the diagonal (from the bottom left corner to the upper right corner) of a heatmap represents the distance of a pharmacy to itself. Dark shades

¹⁷Trottmann et al. (2015), e.g., take the cantons of Aargau and Lucerne for their comparison. However, Aargau is population-wise double the size of Lucerne. Hence, we consider Fribourg the better match.

of grey indicate close distances, in turn lighter shades indicate greater distance. Not only do we see that fewer pharmacies are located in the city centers of Lucerne, the heatmaps also indicate that pharmacies in Lucerne’s city centers are closer to each other and tend to be in the same area.

Figure 5.6 only provide a pair-wise comparison, which is not easy to interpolate, but we can regard this comparison as an example of the long-run impact of different regimes on regional pharmaceutical coverage in allegedly similar city centers. Further, our comparison shows that physician dispensing affects pharmacies in densely populated areas, where additional dispensers are not justified by the original impetus of accounting for scarce population.

Figure 5.6: Comparison of regional coverage in city centers of Fribourg and Lucerne



5.6 Conclusion

In this paper, we exploited Switzerland’s federal legislation to study the effects of different drug dispensation systems on pharmacies within the same country. We

find substantial negative effects of self-dispensation on the utilization of pharmacies as drug dispensers (15% less revenue from drugs). At the same time pharmacies in dispensing regions adapt their portfolios to increase their revenues from non-drugs by 40% compared to pharmacies in non-dispensing regions. These numbers are particularly impressive when we consider that they come on top of massive structural differences.

In 2012, self-dispensing physicians accounted for almost one quarter of the total pharmaceutical market in Switzerland. This influence of physician dispensing becomes apparent in substantial structural differences: cantons that prohibit physician dispensing exhibit more than 33 pharmacies per 100 thousand inhabitants, while self-dispensing cantons only have about 13 pharmacies per 100 thousand inhabitants. Yet, the market volume covered by dispensing physicians and its structural influence shows that, in Switzerland, physician dispensing is clearly established beyond the idea of occasionally filling the coverage gap in scarcely populated areas.

The fact that the total revenue of dispensing physicians exceeds the hypothetical revenue of 730 average pharmacies, which would be necessary to level the coverage gap between dispensing and non-dispensing regions, may indicate topographic differences between regions. However, because both regional groups are dominated by agglomerations and city centers, the additional revenues of physicians can be regarded as a consequence of over-prescription and higher total health care expenditures as suggested in Kaiser and Schmid (2014). Still, because the differences we find are moderate, total health expenditures should not be the only concern of policy makers.

In general, replacement of pharmacies by dispensing physicians affects the accessibility of drugs: First, opening hours of physicians are more limited compared to pharmacies. Second, physicians typically sell to their group of registered patients. Moreover, a physician's portfolio of drugs in stock generally depends on the field of the respective physician. In this context, we also need to consider that the substitution mechanism may be very selective. According to our results, compensation of missing drug revenues by sale of non-drugs is an important strategy of pharmacies in dispensing regions. Sale of non-drugs is driven by walk-in customers and may thus

not be equally successful in less populated areas. In turn, pharmacies may vanish and leave certain regions entirely dependent on local physicians as drug dispensers.

But most importantly, the pharmaceutical market is subject to strict price and margin regulations. Policy makers need to adapt regulations on a regular basis, reacting to budget constraints and dynamic developments in the market. Entertaining a two-tier dispensation system comes at the cost of additional uncertainty about how changes in regulation are absorbed across tiers. In order to avoid undesired effects on total expenditures, accessibility of drugs or coverage with pharmaceutical services, it is important to monitor the effects of legislation on pharmacies closely. For future research, we suggest examining the heterogeneity of pharmacies in greater detail in order to better assess the regulatory effects on different types of pharmacies.

Chapter 6

Technical Aspects: A Software Package to Manage And Archive Economic Time Series Data¹

¹This chapter is based on KOF Working Paper 384 and KOF Working Paper 326.

6.1 Introduction

Time series data are used in a plethora of research fields from meteorology to econometrics. In fields that typically make use of time-varying information of rather low frequency such as monthly, quarterly or yearly data, as opposed to fields that use information obtained from high frequency measuring devices, recording information naturally implies the need for longterm data management. Managing, archiving and providing data consistently over decades can be a challenging task for institutions, which do not have a strong background in software development and data management. In addition to mere technical challenges, sharing social or economic data in modern reproducible research contexts asks for comprehensive context aware meta information since micro-level data can often not be shared at the micro level due to privacy concerns.

In this chapter of the dissertation, I address this challenge and suggest a schema to store time series and their meta information in the proven relational database management system `PostgreSQL`. Alongside the suggested database schema, this chapter introduces the `timeseriesdb` (Bannert, 2015) R package, which adds an additional layer in order to seamlessly access the database from a language that is popular among academic researchers across many fields of empirical research. With the `timeseriesdb` package R, users can access local and remote databases alike, allowing institutions to avoid storing time series information data redundantly in numerous files spread across local hard drives. Using the suggested structure and R package helps managing access using `PostgreSQL`'s role based access management, reduce search costs, while at the same time allowing researchers to conveniently use a familiar programming language.

The idea of providing a software package that is durable, open source, and completely free of license costs is crucial to the work presented in this chapter: Data is stored in `PostgreSQL`, which is widely regarded as the most advanced open source database management system and has been well established for almost 20 years. Further, the unparalleled way the R Language for Statistical Computing and the accompanying environment that has been developed over the last decade has made R an almost natural choice for the open source operational layer of the

`timeseriesdb` project. Today, the R community dedicates an own CRAN Task View (CTV) to time series.² Yet, most time series related R packages are designed to handle times and dates, seasonality, stationarity, unit roots, cointegration, define time series classes, do forecasting and modelling, frequency analysis, decomposition and filtering or resampling. Packages that specifically aim at archiving time series are rather scarce. Thus, `timeseriesdb` aims to add a missing piece to the immense functionality of R by providing the tools to store time series and its (multi-lingual) meta information in a consistent and maintainable way.

More specifically, `timeseriesdb` uses the common database interface R package DBI (Wickham and Müller, 2014) to connect to the database. A connection to directly query the database can then be set up using a database management system (DBMS) specific package, such as `Rpostgresql` (Conway et al., 2013) or `ROracle` (Denis Mukhin and Luciani, 2014). However, these basic interfaces do not provide higher level functionality in the sense that they map time series representations of the operational layer to database tables and vice versa. Rather, these packages can be used to write SQL queries as character strings and send these as queries to the database. Typically, results are returned as standard R *data.frames*. In addition, the pioneering `TSdbi` (Gilbert, 2013a) package along with a family of corresponding DBMS specific packages such as `TSMysql` (Gilbert, 2013b) or `TSPostgreSQL` (Gilbert, 2013c) have addressed the need to conveniently map R time series objects to relations in a database.

Yet, the `timeseriesdb` package distinguishes itself from these existing packages in two specific factors: First, an entire time series is stored in a single table cell as opposed to storing one record per observation. This reduces the number of records substantially, particularly for long time series. Second, `timeseriesdb` enables users to store extensive meta information in multiple languages. Further, the amount

²CTVs monitor, describe and summarize R packages in a particular field. The CTV for time series analysis is maintained by Rob J. Hyndman and can be found at <http://cran.r-project.org/web/views/TimeSeries.html>.

of translated meta information items may vary from record to record. The `time-seriesdb` packages strives to optimize query time when reading from the database or writing to it.

The subsequent section 2 continues to further motivate the use `timeseriesdb` particularly in context of reproducible research and archiving data alongside corresponding meta data. The remainder of this paper is structured as follows: section 3 covers data storage in greater detail and elaborates on the relational structure of the underlying PostgreSQL database. Section 4 gives an overview of the mapping functionality of `timeseriesdb` and shows how R objects are mapped to database relations. Section 5 provides applied code examples to illustrate basic features such as reading and writing from the database. Besides a web-based graphical user interface which ships with the package is introduced. Finally the paper is concluded and an outlook to future releases is given in section 6.

6.2 Motivation

Originally designed to archive economic survey data on the aggregated level, `time-seriesdb` is well-suited to handle any kind of official statistics time series with a daily, monthly, quarterly or yearly release cycle. Establishment statistics are often shared on the aggregated level only, because micro level data – i.e., company or household level data is sensitive. If data is not shared at the micro level and time series cannot be reproduced by the recipient, proper meta information becomes particularly crucial. Against the background of the ongoing development in reproducible research (Koenker and Zeileis, 2009), the ability to trace data back to its provider and describe data extensively becomes a central aspect of an empirical researcher’s work.³ McCullough and Vinod (2003) check five years of American Economic Review articles that involve non-linear solvers for their authors’ efforts

³Buckheit and Donoho (1995) date the origin of the reproducible research movement back to Stanford geophysicist Jon Clearbout, quoting his claim that the scientific publication is not the scholarship but rather its advertising. Consequently, he claims that software is the actual scholarship in a computing-driven research project.

to verify solutions provided by their respective software package. Despite the fact that previous research (Stokes, 2004) had pointed out issues in the reproduction of even top contributions of the profession that involve non-linear solvers, McCullough and Vinod (2003) do not find a single author who reports some kind of efforts to verify results produced by statistical software that involves non-linear solvers. Consequently McCullough and Vinod (2003) claim that publication policies related to the manuscript alone are not sufficient and only both, data and code archives can ensure future reproducibility of empirical economic research. When reproduction problems occur at a later stage, solid code and data archives are mandatory to trace back results and circumstances of the data generating process.

McCullough (2009) shows an overview of economic journals that have made data and code archives mandatory since 1990. By today, the list includes the top journals of the economic profession such as the American Economic Review (AER) which introduced their archive ten years ago in 2004. More recently the idea of reproducible research has experienced yet another push with the R community at the forefront: particularly the `knitr` package (Xie, 2015), (Xie, 2013), (Xie, 2014) and `rmarkdown` (Allaire et al., 2014) are widely used to create documents dynamically from text, code and data. Gandrud (2013) provides a summary of how to create dynamic documents with R including application examples. However, with a reproducible approach to generate reports and research papers data descriptions can be loaded dynamically if comprehensive meta information is available. In the context of archiving data and reproducing research documents Leeper (2014) expresses the need for services that offer data with sophisticated meta data support. Hence, the motivation to create `timeseriesdb` came largely from the intention to design a time series archive that allows storing comprehensive meta information and make this information available in the context of a dynamic process as is the case with `knitr` or `rmarkdown`. Further, multi-lingual meta information should help finding data and thus foster more global exchange of time series information. This meta information becomes important when storing time series information with multiple versions (vintages). National Accounts and GDP time series, for example, are subject to revisions (Shrestha and Marini, 2013), (Branchi et al., 2007) and thus it is important to know which version of time

series was used in the research or forecasting exercise. The ability of `timeseriesdb` to store extensive meta information and make it available in the context of the computational process helps to avoid confusion of series as users can label and select series dynamically. With its license cost free open source components, light weight architecture and low maintenance cost `timeseriesdb` is designed to also suit smaller data providers.⁴

Another important aspect that motivated the development of `timeseriesdb` is its explicit focus on archiving. While other approaches, such as `druid` (Yang, Tschetter, Léauté, Ray, Merlino, and Ganguli, Yang et al.), often execute adhoc computation at the database level at query time, `timeseriesdb` leaves time series operations to R and packages that that have been explicitly designed to process time series. Time series in official statistics can often be the result of complex or computationally intensive processes such as hierarchical aggregation schemes or seasonal adjustment which use specific software. Seasonal adjustment, for example, is most often done with the U.S. Census Bureau’s X13-ARIMA-SEATS Fortran program. The `seasonal` package (Sax, 2014) uses the Census Bureau’s software from within R. Furthermore, R can read numerous foreign file formats (e.g. using `foreign` (R Core Team, 2015) or `R.matlab` (Bengtsson, 2015)) and offers interfaces to various standard software packages as well as to domain-specific software. Thus, R can be used as a flexible interface between other software and the archive database suggested by `timeseriesdb`. In other words, `timeseriesdb` explicitly focuses on storing pre-computed time series and providing a data catalog, while avoiding the need for comprehensive SQL knowledge from the user.

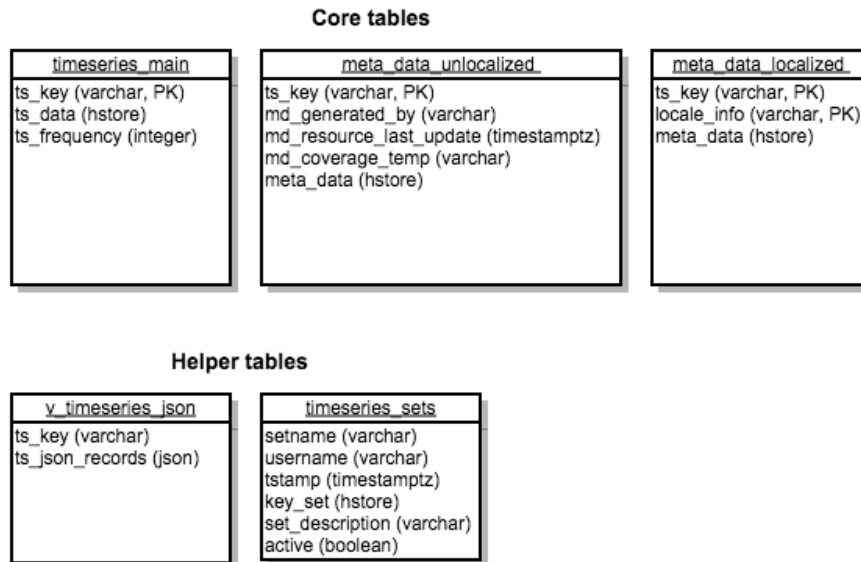
6.3 Data storage

This section covers the relational structure of the PostgreSQL database that is used in `timeseriesdb`. Figure 6.1 provides an overview of all tables and views used in

⁴Please find installation and setup instruction for the database in the appendix to this chapter at the end of the thesis.

the database schema. The following subchapters explain the details of storing time series records as well as corresponding meta information.

Figure 6.1: Overview: relational structure



6.3.1 Data: time series records

The time series records are stored in a simple three column table called *timeseries_main*. While the unique identifier of the record, as well as the frequency of the time series, are stored as standard *varchar* and *integer* data types, respectively, the time series are stored in the PostgreSQL specific *hstore* format. The *hstore* data type stores key-value pairs in a single table cell. In the case of time series, the date is regarded as the key, while the actual value represents the value of the pair. By using *hstore*, an entire time series can be stored in a single row and does not need to be stored in an one-row-per-observation format. PostgreSQL provides functionality to access *hstore* keys inside a row to extract specific parts of the series.

In order to facilitate access, `timeseriesdb` provides a view⁵ called `v_timeseries_json`, which casts data into the popular, more standard *JSON*⁶ format. *JSON* allows for nested structures and thus an entire time series record (key, frequency and time series data itself) can be cast into a single *JSON* entry. With the help of this *JSON*-based view, reading time series into R could be sped up substantially compared to using multiple select statements or splitting results of a single call at the R level.⁷ At the R level `timeseriesdb` makes use of `RJSONIO` (Lang, 2014) to process *JSON*.⁸

6.3.2 Meta data: localized and unlocalized information

While the *hstore* format implements non-relational concepts inside a relational database, meta information is stored in separate relations and linked to the main records. `timeseriesdb` uses two tables to store meta information:

the Table `meta_information_unlocalized` is structured to store meta information that is not translated such as usernames or time span of coverage. The second meta information table, holds meta information that could be translated. Information, such as wording of questions in survey-based time series or elaborate descriptions are stored in the latter table. Again, the *hstore* format is used in both tables to allow for a flexible amount of meta information for each record. For example, some record may have a meta description in French, German and English, while other records only contain German and English meta information. Using *hstore*, this situation can be covered in a single table without storing empty values.

6.3.3 Storing sets of time series

The second helper table, `timeseries_sets`, allows users to store a set of time series keys under a setname (varchar). This can be helpful when users return regularly to retrieve updates of the same series at a later stage. The Table `timeseries_sets`

⁵Views are basically stored select statements in SQL.

⁶JavaScript Object Notation.

⁷see Appendix 'benchmarks' for a detailed summary on the reading speed.

⁸Future versions of `timeseriesdb` will switch to the more recent and optimized `jsonlite` package (Ooms, 2014a) for performance and stability reasons.

will also store the username as a *varchar* and the current timestamp at the time of storage. The set of keys itself is stored in the *hstore* format with the time series key being the key and the type of key being the value. This allows the use of other keys than the primary time series key (e.g. meta information) to identify the time series that belong to a set. In addition a *set_description* can be added as *varchar* and the active flag can be used to activate or deactivate an existing set.

6.3.4 Mapping SQL relations to R objects

The focal functionality of `timeseriesdb` is to map R objects to database relations and vice versa. In general, time series are identified by a unique primary key. Hence, this chapter introduces the most important functions to perform the object - relational mapping between R and PostgreSQL based on identification by a unique time series key:⁹

```
readTimeSeries()  
readMetaInformation()  
storeTimeSeries()  
storeMetaInformation()  
createMetaDataHandle()  
createHstore()
```

6.3.5 Function overview

As the name suggests, the first function reads the time series itself from the database by its key and returns a list that contains at least one R object of class *ts*. The function *readMetaInformation* reads meta information from the database given unique identification of time series and returns an environment that contains meta information. This environment holds meta information objects named like the corresponding time series records. By using a separate environment, objects can have

⁹For more detailed information see the code examples in section 6.4.

the same name as the corresponding time series and are thus easy to link. Further, the resulting meta information environment can be updated by further calls of *readMetaInformation* at a later stage.

Similarly, *storeTimeSeries* stores time series that reside in an R environment or list to the database. Again, a character vector containing time series names can be passed to *storeTimeSeries*. The function *storeMetaInformation* works analogously to the *readMetaInformation*, as it reads meta information from a specified environment and stores them to the database.

The function *createMetaDataHandle* is less essential, but provides an easy way to create convenience operators. The operators help to look up the database using regular expressions. The code chunk below provides a general example for a pre-defined operator that works with the main tables fixed primary key *ts_key*.

```
con %k% 'ts[1-9]{1}$'
```

The object *con* is a PostgreSQL connection object and character string is a regular expression pattern. In this example, the pattern matches all keys $\in ts_1, \dots, ts_9$. The *%k%* operator searches the main table's keys for all records that fit the expression and returns a list of time series that match the query. Such an operator is easy to pre-define because the primary key is fixed and typically not changed by the user. However, querying meta information is very different because users define the amount of meta information keys and their names. Thus, *timeseriesdb* provides a function to create such operators for flexible keys.

```
"%lk%" <- createMetaDataHandle("legacy_key")
con %lk% 'old_key[1-9]{1}$'
```

The example above assumes that there is a meta information field for legacy keys. This means that at least one *hstore* record contains a key called *legacy_key*. After the corresponding operator is created, the operator can be used just like the operator in the fixed key example presented above. Note that the operators created by *createMetaDataHandle* reside in the global environment, as opposed to the namespace of the package and are thus removed when the global environment is

cleared, e.g. by `rm(list=ls())`. The last function in the list, namely `createHstore`, is a method typically not used by the user but is a helper function that is crucial and often used in the mapping process. The function uses the `PostgreSQL` function `hstore` to create key value pairs from different types of R objects. As of version v0.2.1 of `timeseriesdb` methods for classes `ts`, `data.frame` and `list` exists. It is important to understand that `createHstore` does not understand nested formats. Hence, the input is limited to simple, unnested `data.frames` and `lists`.

6.4 Code examples

The following examples illustrate the basic functionality step by step and continue to describe convenience functions, such as plotting a list of time series. For the subsequent examples to work, a database connection to a database that contains a schema called `timeseries` with the relations suggested by `timeseriesdb` is needed. SQL statements for a first time set up of all necessary relations, functions and triggers can be found in the `inst/sql` folder of `timeseriesdb`.¹⁰

The database connectivity of `timeseriesdb` depends on the R packages `DBI` (Wickham and Müller, 2014) and `RPostgreSQL` (Conway et al., 2013).¹¹ The `timeseriesdb` package comes with a convenience function `createConObj` to create a `PostgreSQL` connection object. The `createConObj` function expects several connection parameters, which can be either directly passed to the function or stored as constants using `Sys.setenv`. The advantage of using the system environment as opposed to any object in R's global environment is that it is not affected by clearance of session memory. If the database connection does not change regularly it can be attractive to store the host name, database name and schema either in a global or user-specific `.Renviron` file. User `.Renviron` files are typically located in the user's home directory.

¹⁰Further installation notes can be found in the appendix of this chapter at the end of the thesis.

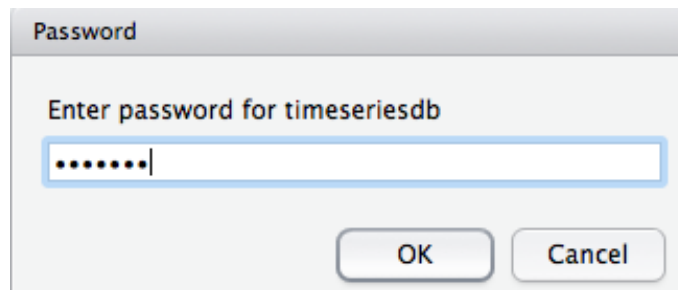
¹¹We will closely monitor the development of `RPostgres` which is currently being developed by Hadley Wickham, but is not an official CRAN package yet. It also relies on `DBI` but has a different design approach. The `Rpostgres` packages might be an alternative to implement database connectivity or could also be supported, additionally.

```
library(timeseriesdb)

# set database name
Sys.setenv(TIMESERIESDB_NAME = "sandbox")
Sys.setenv(TIMESERIESDB_HOST = "localhost")
Sys.setenv(TIMESERIESDB_SCHEMA = "timeseries")
con <- createConObj(passwd = "")
```

The local example database used in this paper does not use a password. Obviously this is not suitable for most production use cases. If password protection is desired, entering passwords interactively should be preferred over storing password in text files. *R Studio* provides the opportunity to use its interface to make use of an interactive password prompt that hides the password from the screen. Figure 6.2 shows *R Studio*'s password dialogue. Also note that the user is taken from *Sys.info* and needs to be specified separately as an argument of *createConObj* if the database user deviates from the system user.

Figure 6.2: R Studio password prompt



6.4.1 Write time series to database

In order to demonstrate writing to the database, a set of 100 random time series is created first. The created time series are of R's standard time series class *ts*. All time series are generated using a random normal and the resulting series are stored in a list. List elements are named with a unique, sequential identifier.

```
# set a seed to make
# the result fully reproducible
set.seed(123)
nms <- paste0("ts",1:100)
ts_list <- lapply(nms, function(x) ts(rnorm(50),
                                     start = c(1998,1),
                                     frequency = 4))

class(ts_list[[1]])

## [1] "ts"

names(ts_list) <- nms
```

The function to store R time series objects to the database is *storeTimeSeries*. It takes names of time series as arguments and searches for corresponding time series in an environment or list. For performance reasons, the function *storeTimeSeries* has been optimized to only use one single SELECT statement to store multiple series. Thus the user should avoid looping over *storeTimeSeries* respectively not use *apply* family functions with *storeTimeSeries*. Rather, a vector of time series names should be passed to the *storeTimeSeries* when multiple series should be stored.

```
storeTimeSeries(nms,con,li = ts_list)

## [1] "100 data and meta data records written successfully."
```

As the output from *storeTimeSeries* shows, information is stored to the main time series table as well as to the meta information table. When a time series

record is stored, a minimal amount of meta data is stored to the unlocalized meta information table. These data are user information, time and timespan covered by the time series. All of this information can be derived from the storage process or the time series itself. Figure 6.3 shows a PostgreSQL client window with a query on the main time series table displaying one of the test time series. Figure 6.4 shows the corresponding minimal meta information that is generated with the initial storage process.

Figure 6.3: PostgreSQL client: a time series record

```

ts_key      | ts1
ts_data     | "1998-01-01"=>"0.922301059107954", "1998-04-01"=>"-2.45149100964402", "199
8-07-01"=>"-0.131003819224936", "1998-10-01"=>"-1.05339701036674", "1999-01-01"=>"1.12716
590131421", "1999-04-01"=>"-0.727834639228732", "1999-07-01"=>"0.935340589208178", "1999-
10-01"=>"-0.468292103835751", "2000-01-01"=>"0.12982106876817", "2000-04-01"=>"1.46235284
119442", "2000-07-01"=>"-0.682169375643072", "2000-10-01"=>"1.81861838674073", "2001-01-0
1"=>"0.986158368247289", "2001-04-01"=>"1.28460132176685", "2001-07-01"=>"-2.246400569400
27", "2001-10-01"=>"-0.168516631588257", "2002-01-01"=>"-1.46661662751573", "2002-04-01"=
>"0.759275035989203", "2002-07-01"=>"1.22277702922995", "2002-10-01"=>"-0.617535389141603
", "2003-01-01"=>"-0.511773942223519", "2003-04-01"=>"-1.62158019298443", "2003-07-01"=>"
0.790937635334595", "2003-10-01"=>"1.46152195521853", "2004-01-01"=>"-1.69993221571343",
"2004-04-01"=>"-1.81251474524571", "2004-07-01"=>"1.14414109742141", "2004-10-01"=>"1.348
54185723644", "2005-01-01"=>"0.371556455073257", "2005-04-01"=>"0.242249032302166", "2005
-07-01"=>"-0.621258548098154", "2005-10-01"=>"0.339038070553804", "2006-01-01"=>"-0.45214
0126031672", "2006-04-01"=>"2.04323321252589", "2006-07-01"=>"-0.449337685075739", "2006-
10-01"=>"-3.13738452799538", "2007-01-01"=>"0.499962209076086", "2007-04-01"=>"-1.2571415
8968218", "2007-07-01"=>"0.822761428159943", "2007-10-01"=>"-1.54609608135241", "2008-01-
01"=>"-0.258780758988273", "2008-04-01"=>"0.390407377124082", "2008-07-01"=>"-0.197270203
079809", "2008-10-01"=>"-1.94694947947191", "2009-01-01"=>"-1.42763817341705", "2009-04-0
1"=>"-0.850418042257596", "2009-07-01"=>"1.62446909297786", "2009-10-01"=>"-0.12663816139
8654", "2010-01-01"=>"1.27560202747805", "2010-04-01"=>"0.179496176715821"
ts_frequency | 4
sandbox=#

```

Figure 6.4: PostgreSQL client: a meta information record

```

sandbox=# SELECT * FROM meta_data_unlocalized WHERE ts_key = 'ts1';
-[ RECORD 1 ]-----+-----
ts_key          | ts1
md_generated_by | mbannert
md_resource_last_update | 2015-03-07 14:22:20+01
md_coverage_temp | 1998-01-01 to 2010-04-01
meta_data       |
sandbox=#

```

6.4.2 Read time series from the database

The corresponding function to read time series from the database is called *readTimeSeries*. Analogous, to *storeTimeSeries*, *readTimeSeries* takes a vector of time series keys and a connection object as its minimal arguments.¹²

```
# clear the memory
rm(list=ls())
con <- createConObj(passwd = "")
results <- readTimeSeries(c("ts1","ts2","ts3"),con)
# note the class of results
class(results)

## [1] "list"    "tslist"
```

Note that *readTimeSeries* always returns a list of time series, no matter how many elements different from zero (i.e. time series) are returned. Again, *readTimeSeries* has been optimized for bulk loading a large amount of time series from the database. Hence, multiple SELECT statements are avoided by the function as it allows processing a vector of time series keys directly. In turn, the user should not use *readTimeSeries* in loops or *apply* constructions.¹³

6.4.3 Adding and reading elaborate meta information

Meta information can either be added at the database level or using R. When adding comprehensive meta information, `timeseriesdb` distinguishes between meta information that cannot be translated (e.g., username or timestamps) and meta information that can (e.g., a written meta description). In both cases, `timeseriesdb` uses the flexible *hstore* format. The use of *hstore* enables `timeseriesdb` to store a different amount of meta information for every record without having to create

¹²This document has been created using the `knitr` (Xie, 2015) package for reproducible research. All code chunks in this paper are executed during compilation of the document. Thus, objects are removed from R's session memory to truly illustrate interaction with the database.

¹³For a more detailed insight on query performance see Appendix: Query Benchmarks.

empty data cells in a rectangular data format for those records that do not have a particular type of meta information.

While a minimal amount of meta information has been stored when a time series record itself was stored, users can add further localized and unlocalized meta information to the database. First, the code chunk below shows how to add additional meta information without a specific locale. The basic idea is to generate a separate environment to hold all unlocalized meta information while the time series may reside in the *.GlobalEnv* or any other environment distinct from the meta environment. Meta information is meant to have the same object name as the corresponding time series.

```
# add seed info that was
# used to create the series and some legacy id
meta_ts1 <- list(seed = 123,legacy_key = 'series1')
meta_ts2 <- list(seed = 123,legacy_key = 'series2')

meta_unlocalized <- addMetaInformation("ts1",meta_ts1)
addMetaInformation("ts2",meta_ts2,meta_unlocalized)

## x contains 2 meta information object(s).
```

The same function can be used in an analogous fashion to create an environment that holds localized meta information. The following example adds English meta information to the time series *ts1*. Note that the locale argument of *addMetaInformation* defaults to NULL and thus only needs to be set if and when localized meta information is added.

```
en <- list('short_description' = 'Random Series',
          'full_description' = 'Random Normal using seed 123.')
```

```
meta_en <- addMetaInformation('ts1',en)
meta_en$ts1
```

```
## $short_description
## [1] "Random Series"
##
## $full_description
## [1] "Random Normal using seed 123."
##
## attr(,"class")
## [1] "miro" "list"
```

In both cases meta information is stored in separate R environments. In a second step we can store all records contained in these environments to the database as shown in the following code chunk.

```
# store localized Information
storeMetaInformation("ts1",con,locale = "en",
                    lookup_env = "meta_en",overwrite = T)

## [1] "en meta information successfully written."

# store unlocalized information
storeMetaInformation("ts1",con,tbl ="meta_data_unlocalized",
                    "meta_unlocalized",
                    locale = NULL,
                    overwrite = T)

## [1] " meta information successfully written."
```

Similarly, meta information can be read from the database into an R session. Loading meta information is done explicitly because meta information should not be implicitly loaded into the R session when reading in time series. Also, users should be able to vary the amount of meta information that is loaded into context. Depending on the different use cases, such as batch processing or exploratory data research, a different amount of meta information may be desired. Thus, reading meta information is separated from reading data.

```

m_unlocal <- readMetaInformation('ts1',con,NULL,
  tbl = "meta_data_unlocalized")
m_local <- readMetaInformation('ts1',con,"en")
m_unlocal$ts1

## $fixed
##   md_generated_by md_resource_last_update      md_coverage_temp
## 1      mbannert      2016-11-22 14:59:28 1998-01-01 to 2010-04-01
##
## $flexible
##      key    value
## 1     seed    123
## 2 legacy_key series1
##
## attr(,"class")
## [1] "miro" "list"

m_local$ts1

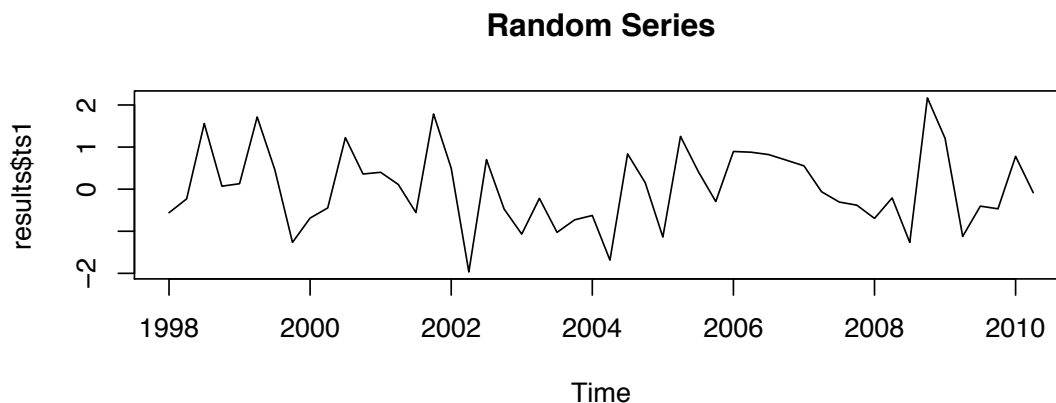
## $full_description
## [1] "Random Normal using seed 123."
##
## $short_description
## [1] "Random Series"
##
## attr(,"class")
## [1] "miro" "list"

```

Meta information can also be used to label figures and tables dynamically. Figure 6.5 shows a time series plot with a dynamically created title loaded from a PostgreSQL database.

Figure 6.5: Time series with dynamic meta description

```
plot(results$ts1, main = m_local$ts1$short_description)
```




6.4.4 Plot a list of time series

The `timeseriesdb` package will return a list of time series to the R session when the database is queried for time series records. Because base R does not provide a function to plot a list of time series out of the box, `timeseriesdb` comes with a convenience method that accepts a list with the additional class 'tlist' containing one or more elements of class "ts". Thus, `timeseriesdb` functions that return lists of time series, namely `readTimeSeries` as well as the quick handle operators presented in Section 4, add the class attribute 'tlist' to the returned list. The user can simply use the plot method shown below to plot all time series contained in such a list.

The dimensions and axes will be automatically created to suit all series. Internally, the canvas is built using the most extreme values out of all time series. Subsequently, further series are added using the basic `lines` command. Figure 6.6 shows a list of time series plotted using the `tlist` method for `plot`.

Figure 6.6: Plot method for list of time series

```
plot(results)
```



6.4.5 Data explorer

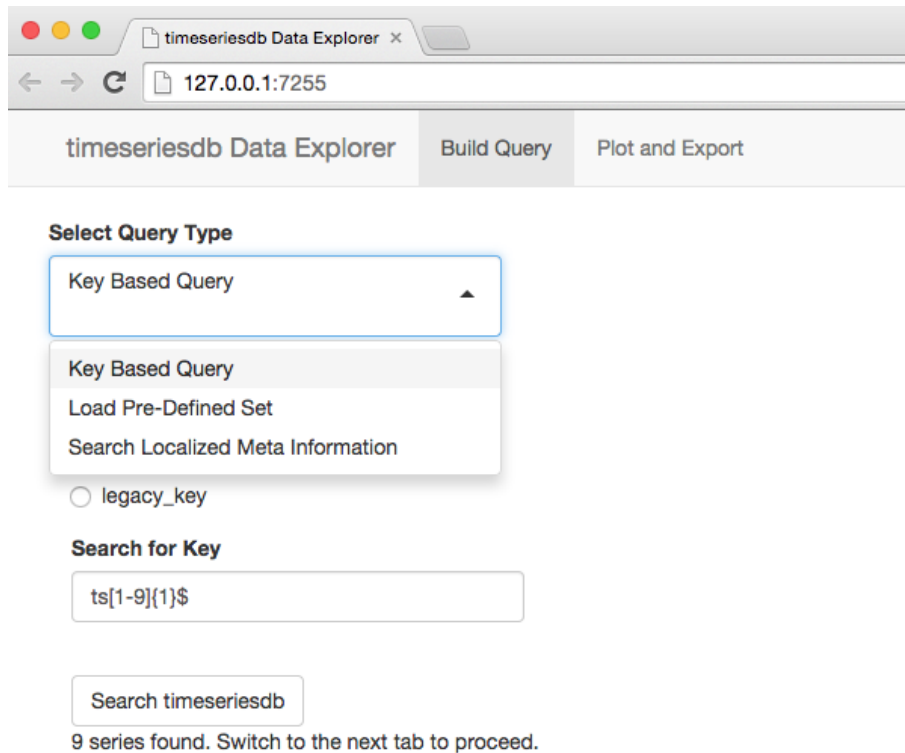
Another way to explore the content of a `timeseriesdb`-based database is the *data explorer*. The data explorer is a simple web-based graphical user interface (GUI) based on the popular `shiny` (Chang et al., 2015). Users can skim data interactively given a database connection. To start the GUI, run the code shown below, where `con` is a valid PostgreSQL connection object.

```
exploreDb(con)
```

Once the GUI has started, the user has three options to query the database using the GUI: A *Key Based Query* lets the user look for matches with either the time series' primary key or within fixed meta information keys. The latter can be a relevant alternative when short cut aliases or legacy keys are used. The second option is to query localized meta information. Third, users can load pre-defined sets

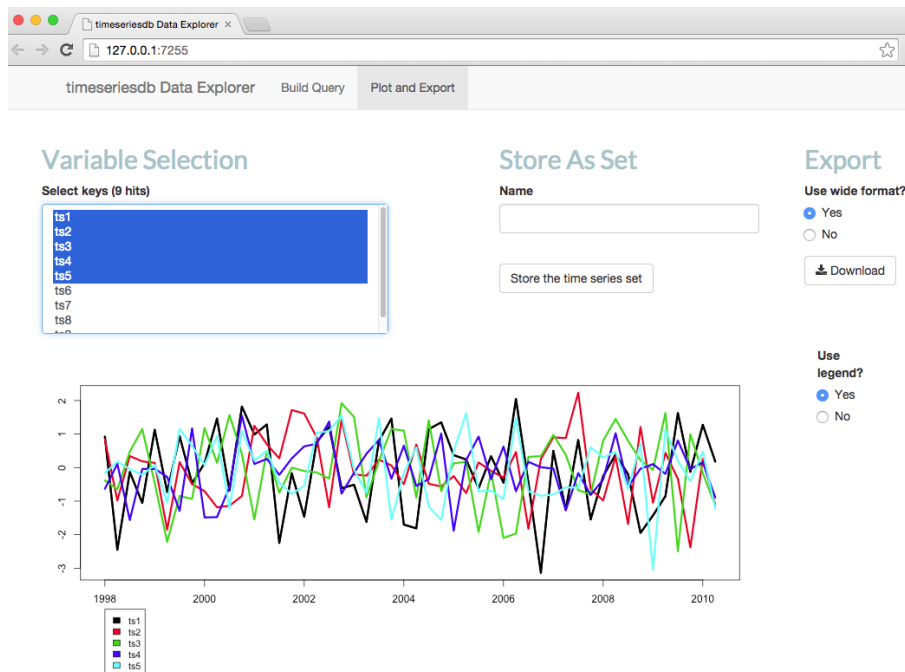
of time series they have stored previously. In any case, all query types return zero or more time series keys. Figure 6.7 shows the tab to build queries.

Figure 6.7: Web based data explorer: build queries



After creating a valid query, the user may switch to the *Plot and Export* tab and select those series that should be plotted, exported or saved in a set. After a selection is made, the keys of the selected series can be stored under a set name. By doing so, the set will be stored to the database and can be loaded again at a later stage. Also, time series can be exported as .csv files in either wide or long format. Selected series will immediately be plotted below to give users an overview of the selected series. Note that R's own localhost, which hosts the shiny-based data explorer, needs to be stopped if the user wants to go back to R's interactive console. Otherwise a separate R session needs to be started. Figure 6.8 shows the *Plot and Export* tab.

Figure 6.8: Web based data explorer: plot and export



6.5 Conclusion and outlook

Although, the brief overview provided in this paper cannot cover all details of `timeseriesdb`'s functionality, the general idea and core functionality have been presented: `timeseriesdb` is a package to manage and archive time series data along with comprehensive meta information. Using `R` and `PostgreSQL`, the closest existing CRAN package is `TSPostgreSQL`. Besides queries optimized for bulk storage and the use of the `hstore` data type, the main distinction of `timeseriesdb` from `TSPostgreSQL` is its support for extensive meta information. Using different storage types and optimizing queries comes at the cost of not being able to flexibly switch the DBMS like the `TSdbi` package family does (e.g. to `MySQL`), but `PostgreSQL`'s advanced and comprehensive approach is worth the loss of this particular flexibility.

Future versions of `timeseriesdb` aim to improve support for writing large sets of time series and split large query strings to avoid `R` memory issues. Also, future versions intend to extend the support for irregular time series, which go hand

in hand with support for Zeileis and Grothendieck (2005) and Ryan and Ulrich (2013) time series representation. Though R related functionality will be extended, `timeseriesdb` is not limited to the *R Language for Statistical Computing*. Because the package uses a standard PostgreSQL database, many software packages, including web technologies, can connect to the database. Exploiting these opportunities will be an important expansion for the package: common web languages could be used to create a REST Application Programming Interface (API) to access and export timeseries data using a standard web browsers. Thanks to advances in the development of technologies that bring R to web servers, extending `timeseriesdb`'s functionality to create outputs that can be used to provide a web service seems generally appealing. Most prominently, `shiny`, which has already been used to create the web GUI inside `timeseriesdb`, could be used to create such a service using *shiny server*. Also, `opencpu` (Ooms, 2014b) provides a promising alternative approach that completely separates concerns and allows users to use functions of their own packages via a web service.

Appendix A

Appendix to Chapter 2

Feedback Questionnaire



KOF KOF Konjunkturforschungsstelle
ETH Zürich, WEC D 19, 8092 Zürich
www.kof.ethz.ch

Tel. 044 632 80 64
Fax 044 632 13 52



Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Please note:
- Confine your answer to the sector specified in the questionnaire
- Consider the business activity of offices in Switzerland only
- Please colour in the appropriate field
- Read the notes overleaf
- Return before 30 June 2013
All data will be treated in the strictest confidence

1. Assessment of the Business Situation

a) Please tick the importance to you of each of the following factors in assessing the current business situation:

| | unimportant 1 | 2 | 3 | 4 | very important 5 |
|---|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| Turnover | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Quantity of sales/volume | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Inventory | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Enquiries/offers | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Incoming orders | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Customer frequency | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Corporate earnings/margin | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Cost situation | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Liquidity | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| General sentiment in the sector/overall economy | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Other: _____ | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

b) Do you use a yardstick to assess the business situation? If so, how do you compare it?

- With the targets/expectations/ for the month under review
- With the situation of major competitors or the situation in the sector
- With the general economic situation
- With a previous business situation
- Other: _____
- We don't use a yardstick

↳ If you make a comparison with a previous business situation: with which business situation do you compare it?

- Previous month
- Same month of the previous year
- Same month of the previous year
- Average business situation in recent years
- Business situation in the corresponding month of recent years
- Other: _____

2. Business Expectations

Please tick the importance to you of the following list of factors in assessing the business performance expected in the (coming 6 months):

| | unimportant 1 | 2 | 3 | 4 | very important 5 |
|---|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| Factors/indicators of own business | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Economic forecasts for our sector / for the overall economy | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| General sentiment in the sector / overall economy | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| General sentiment in the client sector | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Trends in household earnings /labour market situation | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Macroeconomic domestic production / investment | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Macroeconomic imports / exports | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Political/general legal framework | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

3. Seasonal Fluctuations

We ask our survey participants to eliminate purely seasonal fluctuations from questions on business and employment expectations. If this is feasible, what approach do you use to eliminate purely seasonal fluctuations?

- Elimination of seasonal fluctuations not possible or feasible
- Comparison with previous year
- Comparison with average trends in the same months of previous years
- Empirical values
- Other: _____


4. Business Activity

Could your company's existing capacities cope with increased demand at present?

- Yes No

↳ If yes: We could increase our business activity by % ausweiten.

- Question is irrelevant for our company



5. Number of Surveys

a) How many surveys do you (personally) complete on average in a year?

| | 0 | 1 | 2 | 3 | 4 | >4 |
|-------------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| KOF | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Official statistics/FSO | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Market researchers | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Industry association | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Other: _____ | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

b) Does your company have a routine for replying to surveys?

The KOF Business Tendency Surveys are usually...

...answered immediately after receiving the questionnaire

...answered when time permits

...answered just before the deadline for submission

...answered when a reminder is received

c) How long does it usually take to fill out the KOF Business Tendency Surveys (in minutes)?

| <5 | <10 | <15 | <20 | <30 | <45 | <60 | >60 |
|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

6. Personal Data

a) In which department of your company are KOF Business Tendency Surveys normally completed?

Management

Staff Position

Controlling

Accounts

Statistics

Procurement

Sales

Marketing

Other: _____

b) What position does the person who usually completes the surveys occupy in the company? (With reference to the questionnaire)

Owner/manager/board member

Departmental head

Administrator

Other: _____

c) How many people are involved in completing the KOF Questionnaire (Quarterly Business Tendency Survey)?

one two three more than three

7. Significance of Business Tendency Surveys

How do you rate the benefit of business tendency surveys in general?

| unnecessary | 1 | 2 | 3 | 4 | important | 5 |
|-------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

8. Capacity Utilisation

a) Please tick the importance to you of each of the following factors in assessing the capacity utilisation of your business:

| | unimportant | 1 | 2 | 3 | 4 | very important | 5 |
|---|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| Working hours | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Personnel | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Productivity per employee | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Tech. Capacities | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Surface area/existing stock of buildings and premises | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Material consumption | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Capital used | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Licences | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Customer frequency | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Other: _____ | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

b) If the demand for your services exceeds your existing capacities, would you be able to increase your capacities at short notice?

| | impossible | 1 | 2 | 3 | perfectly possible | 4 | 5 |
|--------------------------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| Personnel-related factors | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Other, non personnel-related factors | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

9. Recording Capacity Utilisation

a) If possible, assess the average capacity utilisation in the last 3 months based on...

...personnel-related capacity (in %):

| 50 | 55 | 60 | 65 | 70 | 75 | 80 | 85 | 90 | 95 | 100 | 105 | 110 | 115 | 120 |
|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

or: question is not relevant to our company

...non personnel-related, technical capacity (in %):

| 50 | 55 | 60 | 65 | 70 | 75 | 80 | 85 | 90 | 95 | 100 | 105 | 110 | 115 | 120 |
|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

or: question is not relevant to our company

b) If the demand for your services exceeds your existing capacities, would you be able to increase your capacities at short notice?

| productive | unproductive | not specified |
|--------------------------|--------------------------|-----------------------|
| <input type="checkbox"/> | <input type="checkbox"/> | <input type="radio"/> |

c) Do you use a time recording system?

Yes

No, but we plan to introduce one

No

10. Staff Shortage

a) Have your activities been hindered in the past by a staff shortage?

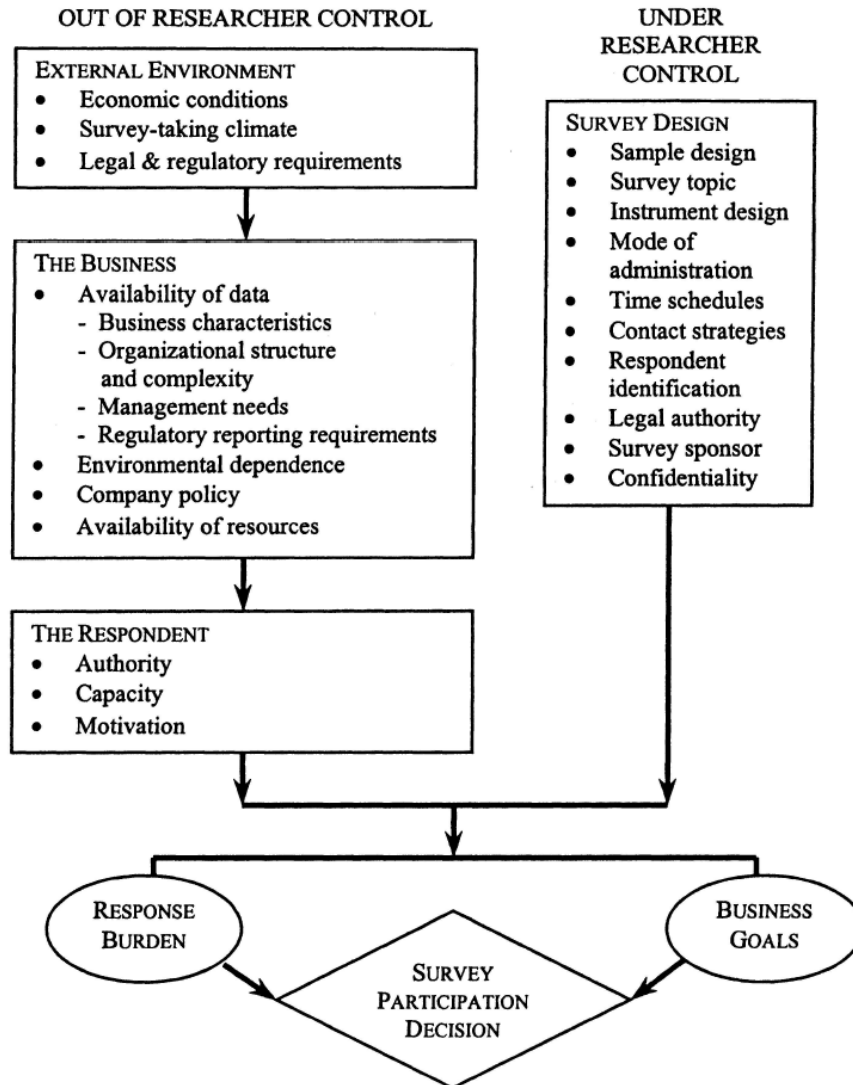
Yes No

b) If there were problems recruiting suitably qualified staff, this was attributable to

the lack of availability of staff with relevant qualifications

the excessive costs of employing staff with relevant qualifications

Conceptual framework for non-response in BTS



Appendix B

Appendix to Chapter 3

Validity of the survey experiment

This section discusses the general challenges that need to be overcome in order to obtain reliable data from surveys. The discussion of survey errors in this appendix covers the different aspects described in the handbooks of Dillman et al. (2014) and Weisberg (2005): sampling error, measurement error due to interviewers or respondents, coverage error, non-response and post-survey error. The respective subsections will discuss these aspects of error and what we have done to avoid them in practice. Interviewer bias will not be discussed in detail as our survey is conducted as an online and paper-based mixed-mode survey only and is therefore not prone to interviewer bias. The only personal contact with participants is through telephone-based reminders that do not interview participants. Further, post-survey error is not discussed in great detail as the survey is conducted by the KOF Swiss Economic Institute, which has several decades of experience in conducting surveys and maintaining databases that store survey data. Thus, possible post-survey errors related to not storing answers correctly or difficulties reading survey data are not discussed at length. In turn, we place additional emphasis on issues that relate to economic surveys, in particular, bias introduced by inaccurate respondents, coverage, and non-response. Further, Druckman et al. (2006), Kinder (2007), Gaines et al. (2006), Barabas and Jerit (2010), and Guterbock and Nock (2010) carefully examine possible pitfalls of survey experiments.

Measurement error due to respondents

Measurement errors due to respondents can be the consequence of several aspects, such as firms' understanding of the questions, relevance of the questions to participants, willingness to answer the questions correctly and carefully or shortcomings of the questionnaire. We checked the validity and relevance of our scenario questions carefully in an interviewer pre-test among randomly selected companies.¹⁴ We also

¹⁴Groves (1989, p. 422) writes that the recall of past events depends on the length of the recall period, the salience of the event to be recalled, the task difficulty of the event, and the respondents

ensured that the questionnaire reached out to the right persons, namely top-level executives or heads of accounting/controlling by attaching the scenario questions to the bi-annual KOF Investment Survey. The KOF Panel is well established among companies and is continuously maintained to keep high-quality relations with Swiss firms in the service, the construction, and the manufacturing sector. Surveys based on the KOF panel have a long history of contributing to valid predictions of various economic figures, and thus, using the KOF sample encourages our trust in firms willingness to participate carefully and correctly.

We do not claim that every single firm executive forms unbiased and informative expectation. Rather, for survey-based impulse responses to be valid the responses must be unbiased in the aggregate (i.e., on average over all firms). This weak form of unbiasedness can be tested by checking whether firms realizations are in line with their expectations in the baseline scenario. In order to do so, we use information from two waves of the KOF Investment Survey, which collects firms' expectations and realizations of key financial figures. The 2012 spring wave, which is linked to our scenario survey contains projections and the 2013 spring wave contains the corresponding realizations for the same period. Thus, we can directly compare firms' projections with their realizations. Because no substantial macroeconomic shock hit the Swiss economy during this period, we can assume that deviations are rather based on idiosyncratic shocks or measurement errors due to the respondent. When testing for differences, we ensure data quality by excluding observations with obvious typos and multiplication/division by 1000 mistakes (i.e., projection 1000 Swiss Francs, realized 1 million Swiss Francs). In order to test if projected numbers are on average in line with realized figures, we conduct a Student t-test on projection errors by calculating the firm-specific relative projection error

$$e_{GI_i} = \frac{GI_i - GI_i^e}{GI_i^e},$$

attention or motivation. More recent events are recalled with more precision than events further in the past. Specifically, Gaines et al. (2006) argue that agents are indeed able to evaluate the causal effects of a counterfactual event provided that they experienced similar events in the past, and that they consider such events to be relevant.

where GI_i is gross investment reported by firm i and GI_i^e represents projected gross investments. Table 6.1 shows the results of a t-test testing the null hypothesis that the true mean is equal to 0. With a p-value of 0.29, it cannot be rejected that the mean of 0.027 is equal to zero. Hence, it cannot be rejected that firm projection errors are i.i.d. and the average projection error is zero.

This finding confirms the ability of our respondents' to consistently predict the outcomes of their companies' financial figures. Consistent expectations in the baseline case also validate the scenario case because the scenario can be considered as a simplification of the baseline (see section on survey validity in chapter 3).

Table 6.1: Student t-test on firms' projection errors

| Mean projection error | t-values | p-values | obs. |
|-----------------------|----------|----------|------|
| 0.0265 | 1.052 | 0.293 | 657 |

Coverage error

Coverage error occurs when there is a bias due to the omission of noncovered units (Weisberg, 2005, p. 206). In our context, this might be the case if we underrepresent certain industry types while overrepresenting others. In order to correct for such misalignments, we give extra weight to hard-to-obtain respondents. The answers are weighted twofold (see aggregation scheme in this appendix). First, individual responses within an industry group (based on NACE classification scheme) are weighted based on firms' number of employees. This weighting scheme yields the distribution of answers for each industry category. Second, the industry categories are aggregated to an industry average by utilizing gross value added shares provided by Swiss Statistical Office. The number of employees is a proxy for the value added of a company. Industrial firms with a larger value added are more likely to also employ more workers. By weighting the answers first for the industry categories and second for industry total, we ensure representativity of our sample. This aggregation scheme is in line with international standards for the aggregation of business sentiment surveys (European Commission, 2007).

Accounting for non-response bias

Potentially our results could be biased through self-selection into the sample. If our questionnaire was highly relevant to a particular group of firms, these firms may systematically select themselves into the sample, while firms for whom the questionnaire is less relevant, might choose to drop out. In that sense, if our questionnaire was only relevant to firms with a large oil share, our result might over-estimate the effect of an oil price shock for the entire economy because highly oil-dependent firms were over-represented by selecting themselves into the sample.

To ensure our estimation results are not biased by the fact that firms with an high oil share were more likely to respond than others, we apply the “surrogate” approach of Wallace and Mellor (1988). We compare the firms that responded on time (i.e., by July 9, 2012) with those that did not answer the survey until that date. Regarding the late responders, we enforced participation by gently pressuring them with phone calls. Hence, the late responders can be interpreted as a sample from the non-response population. Following Wallace and Mellor (1988), we create two subsamples by selecting the first 50 observations from the early responders and the last 50 observations from late responders. Given that subsample participants submitted in random order, both subsamples should be random draws from the total population and thus should not differ in their means and distribution. Based on this idea, we test for differences in the means for the oil share variable. Note that this could be done in principle for differences in any of our variables; however, we focus on the oil share because of its focal role in our scenario. We also perform a Kolmogorov-Smirnov test to check whether distributions of both subsamples are equal. Next we present the results for an unweighted as well as a weighted sample, which weighs firms according to their contribution to the entire economy.

Table 6.2 shows the results of a simple t-test for the equality of means in both groups: late respondents and early respondents. The differences in mean oil shares are not significant at the 10 percent level. This results also holds when we multiply firms’ oil shares by their weights used in the aggregation procedure.


The sample mean provides a first glance at the oil share variable but does not fully represent the differences in sampling distributions. Thus, we also look at the

Table 6.2: Mean oil shares of early and late respondents

| Test | Mean oil share of early respondents | Mean oil share of late respondents | p-value |
|-------------------|-------------------------------------|------------------------------------|---------|
| Unweighted t-test | 4.4107 | 2.6941 | 0.1880 |
| Weighted t-test | 0.0010 | 0.0043 | 0.1050 |

sampling distributions of both samples. Figure 6.9 shows histograms for both groups and both samples.

Figure 6.9: Distributions by samples



The Kolmogorov-Smirnov (KS) tests fails to reject the hypothesis of equal distributions for both the weighted and the unweighted sample. Table 6.3 summarizes the results of the KS test.

Table 6.3: Kolomogorov-Smirnov test results for equality of distributions

| Test | Test statistic D | p-value |
|--------------------|------------------|---------|
| Unweighted KS-test | 0.1980 | 0.4147 |
| Weighted KS-test | 0.2544 | 0.1512 |

Enterprise panel

The enterprise panel of the KOF Swiss Economic Institute is based on a sample of 7000 firms taken from the Business Register (BR) of the Federal Statistical Office. The sample, which covers manufacturing, construction, and the commercial area of the service sector, is stratified according to sectors and sector-specific variables and is adjusted regularly. The corresponding address database, which in addition to numerous structural features of the firms, contains all information that is required to ensure that panel surveys run smoothly (specific contact points, checking of replies, recalls, and incomplete surveys etc.), was developed continuously and updated regularly throughout the project.

Based on the KOF enterprise panel, regular collections of data on the structural changes, innovative activities and investment plans of Swiss companies are conducted. The enterprise panel draws from a population of 60270 companies. These companies are part of the entire collection of firms within Switzerland. Out of all firms within Switzerland only those with a business register number (assigned in year 2001) and with at least 5 employees were selected. Companies belonging to agricultural activities or public administration are also included. In contrast to the sample used for structural investigations of the Swiss economy, the utilized investment sample also contains firms active in education, health, waste disposal, entertainment, cultural and sports activities.

The sample of 7000 firms was drawn from the population of 60270 companies by utilizing stratified random sampling. While simple random sampling would assign the same probability of being drawn to each observation (or firm), stratified random sampling allows adjustments to the sampling weights. Stratification is important in order to achieve representation at the sector and industry level, as well as in terms of size classes. Size classes are proxied by the number of employees. Larger

companies are assigned higher sampling weights compared to smaller companies, firms being active in less represented industries are also assigned higher weights. The population of 60270 firms has been divided into three size classes within industry groups: small, medium, and large, based on the number of employees. Within each industry, the cut-offs between small, medium and large differ. The cut-offs have been determined based on the distribution of firms size within an industry. In the Swiss financial industry and woodworking sector, the smallest banks have more employees than the smallest carpenter and the largest banks have more employees than the largest carpenter. Therefore, each industry requires individual cut-offs, based on the distribution of employees within an industry. These cut-off values are determined according to Dalenius and Gurney (1951).

Once cut-offs of size classes within industries have been determined, the sampling weights within each industry size class are assigned. The sampling weight (i.e., the drawing probability) of large firms within each industry has been set to 100%. All large firms shall be included (this increases the number of employees covered by the survey). The sampling weights of medium categories are a mixture between 100% drawing probability and random sampling. The medium-sized firms, which have been drawn with 100% drawing probability belong to the chemical industry, metal production, machinery, electrical engineering, electronics and instruments, watches, cars, energy, retail, transportation, banking and insurance, and communication. Industries with a priori assumed higher tendency to innovate have been assigned a higher drawing probability in the medium category. Drawing probabilities for the small classes and remaining medium classes are based on the number of firms within each industry size class multiplied by a weighting factor and divided by the sum of all other number of firms again multiplied by the same weighting factor. Thereby, the relative weight of each industry size class (small and medium) not assigned a 100% drawing probability, can be determined. This drawing probability is then multiplied by the number of firms still to be drawn. The number of remaining firms is the difference between the target sample size (i.e. 3000 for manufacturing, 600 for construction, 3400 for services) and the number of firms with 100% drawing probability (i.e., for manufacturing $3000 - 929 = 2071$). Larger are assigned to the

remaining medium-sized firms compared to the weighting factors for small firms. These are determined endogenously based on the number of full-time worker equivalents. Within each industry size class, the sampling weights are the same. The attribution of sampling weights for each industry size class is based on Cochrane (1977).

The sample size of individual industries has been adjusted according to

$$\tilde{n}_i = (N - C) * \frac{n_i s_i}{\sum n_i s_i} \quad (6.1)$$

with \tilde{n}_i the adjusted sample size of different industries in a sector, separated by size, and N the target sample size of a sector. C represents the correction of the target sample size, calculated as the sum of all estimated samples that are not adjusted. n_i represents the size of an industry that has to be adjusted and s_i its standard deviation (Cochrane, 1977, p. 104). The sampling weights w_i are then calculated as

$$w_i = \frac{\tilde{n}_i}{n_i} \quad (6.2)$$

Our sample contains 7000 firms. Out of these, 3000 firms belongs to manufacturing, 3000 belong to services, 600 to construction and 400 to health services, waste disposal, education, cultural activities and sports. Within each of these categories, the distribution of weights is optimal (see Cochrane, 1977).

This questionnaire has been completed by:

In future we would like to answer the questionnaire via the internet. Our e-mail-address:

Name:

.....

Function:

Telephone:

Many thanks for your participation

Explanations

General remarks

The Investment Survey is an instrument for the early recording of planned investment trends.

Definition «Turnover»

The turnover addressed by this questionnaire conforms with the definition of the Swiss Federal Statistical Office:

Definition «Investment»

The investments addressed by this questionnaire mean inflows minus outflows of fixed capital assets. These should be recorded before depreciation on the basis of their purchase price (gross investment). It is irrelevant whether the equipment which is being used for the first time is new or second-hand, and whether it has been bought, hired or created in-house.

«Turnover comprises the totals invoiced by the observation unit during the reference period, and this corresponds to market sales of goods or services supplied to third parties. Turnover includes all duties and taxes on the goods or services invoiced by the unit with the exception of the VAT invoiced by the unit vis-a-vis its customer and other similar deductible taxes directly linked to turnover. Turnover also includes all other charges (transport, packaging, etc.) passed on to the customer, even if these charges are listed separately in the invoice.

Fixed capital formation thus encompasses:

Construction

- New construction, conversion work and renovation of commercial premises.

Reduction in prices, rebates and discounts as well as the value of returned packing must be deducted. Price reductions, rebates and bonuses conceded later to clients, for example at the end of the year, are not taken into account.»

Machinery and equipment

- Machinery, mechanical plants, conveying equipment and warehouse equipment, office machines incl. IT (hardware and software), furniture and equipment, vehicles used for business purposes, and (only) such services which are designed to preserve, to improve or to renovate plants.

Banks:

Earnings from interest revenue and trading, services and commission business.

This means that fixed capital formation does not include:

- Financial investment (e.g. equity holdings, securities)
- Investment in residential property
- Real estate costs
- Buildings and plants which are intended for hire by the lessor, where the lessor acts merely as a (third-party) financier
- Inventory investment (inventory increases)
- Intangible assets (e.g. expenditure on marketing concepts, for human capital, for research & development, for patents and licences)

Insurances:

Gross premiums minus gross payments for insurance claims plus net earnings from capital investments; gross fees for consulting services.

Definition «Expenditures»

Expenditures are defined as expenses for material, goods and services, wages and labor costs, social security contributions, other personnel and operating expenditures.

No expenditures are therefore:

Investments, financial expenses, depreciation, other write-downs, additional costs, nonoperating and extraordinary expenses, taxes.

Remarks

KOF Special Survey

«Oil Price and Exchange Rate»

KOF Konjunkturforschungsstelle
ETH Zürich, WEC D 15, 8092 Zürich
www.kof.ethz.ch

Tel. 044 632 85 33
Fax 044 632 13 52
ivu@kof.ethz.ch

An approximate estimation based on experience is sufficient. Precise figures are not required.

6. Exchange Rate

- a) The Swiss National Bank (SNB) announced to defend the lower limit of 1.20 CHF/EUR. The current exchange rate of Euro to Swiss Franc is 1.20 CHF/EUR. Which average exchange rate do you expect?

2nd half of 2012
 ≤1.00 1.10 1.15 1.20 1.25 1.30 ≥1.40 NA

2013
 ≤1.00 1.10 1.15 1.20 1.25 1.30 ≥1.40 NA

- b) How large are your current exports as a percentage of total turnover?

Exports to Euro Area
 0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100% NA

Exports to Rest of the World
 0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100% NA

- c) How large are your current imports as a percentage of total turnover?

Imports from Euro Area
 0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100% NA

Imports from Rest of the World
 0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100% NA

7. Scenario «Exchange Rate»

Suppose the SNB reduces the lower limit of the exchange rate to 1.10 CHF/EUR under else constant economic circumstances. Suppose this leads to an exchange rate of 1.10 CHF/EUR, which corresponds to a revaluation of the Swiss franc. How your financial figures change under these circumstances as compared to your previous expectations for these figures?

- a) Total Turnover

2nd half of 2012
 ≤-7.5% 5% -3% -2% -1% 0% 1% 2% 3% 5% ≥7.5% NA

2013
 ≤-7.5% 5% -3% -2% -1% 0% 1% 2% 3% 5% ≥7.5% NA

- b) Total Expenditures (incl. staff, inputs, other expenses; excl. investments)

2nd half of 2012
 ≤-7.5% 5% -3% -2% -1% 0% 1% 2% 3% 5% ≥7.5% NA

2013
 ≤-7.5% 5% -3% -2% -1% 0% 1% 2% 3% 5% ≥7.5% NA

8. Oil Price

How large are your expenses for oil (e.g. fuel, gasoline, diesel, oils, grease, plastics, chemical products) as a percentage of total expenditures?

0% 1% 2% 3% 4% 5% 7.5% 10% 12.5% 15% 20% ≥30% NA

9. Scenario «Oil Price»

Suppose the oil price increases by 30% within the next month under else constant economic circumstances and will remain 30% above your previous expectations regarding the oil price development. How do your financial figures change compared to your previous expectations regarding these figures?

- a) Purchase Prices (average of all purchases of goods and services)

2nd half of 2012
 ≤-7.5% 5% -3% -2% -1% 0% 1% 2% 3% 5% ≥7.5% NA

2013
 ≤-7.5% 5% -3% -2% -1% 0% 1% 2% 3% 5% ≥7.5% NA

- b) Total Expenditures (incl. staff, inputs, other expenses; excl. investments)

2nd half of 2012
 ≤-7.5% 5% -3% -2% -1% 0% 1% 2% 3% 5% ≥7.5% NA

2013
 ≤-7.5% 5% -3% -2% -1% 0% 1% 2% 3% 5% ≥7.5% NA

- c) Domestic Sales Prices

2nd half of 2012
 ≤-7.5% 5% -3% -2% -1% 0% 1% 2% 3% 5% ≥7.5% NA

2013
 ≤-7.5% 5% -3% -2% -1% 0% 1% 2% 3% 5% ≥7.5% NA

- d) Foreign Sales Prices

2nd half of 2012
 ≤-7.5% 5% -3% -2% -1% 0% 1% 2% 3% 5% ≥7.5% NA

2013
 ≤-7.5% 5% -3% -2% -1% 0% 1% 2% 3% 5% ≥7.5% NA

- e) Total Turnover

2nd half of 2012
 ≤-7.5% 5% -3% -2% -1% 0% 1% 2% 3% 5% ≥7.5% NA

2013
 ≤-7.5% 5% -3% -2% -1% 0% 1% 2% 3% 5% ≥7.5% NA

Data

Table 6.4: Variables (KOF Survey and external sources)

| Variable | Description | Scale |
|--|---|--------------------------------|
| Oil Share | Share of expenses for oil (e.g. fuel, gasoline, diesel, oils, grease, plastics, chemical products) as a percentage of total expenditures | 0% to $\geq +30\%$ |
| Export Share Euro Area | The share of exports to Euro area countries relative to total turnover | 0% to +100% |
| Export Share World (excluding Euro Area) | The share of exports to countries outside the Euro area | 0% to +100% |
| Import Share Euro Area | The share of imports from Euro area countries relative to total cost | 0% to +100% |
| Import Share World (excluding Euro Area) | The share of imports from countries outside the Euro area | 0% to +100% |
| Employees | Number of employees (full time equivalents) in Switzerland at the end of 2011 | Absolute values |
| Turnover Nominal | Nominal turnover excl. VAT generated by the Swiss parts of the company (including sales to foreign countries) in Swiss Francs. Reported balance sheet values for 2010 and 2011, projected values for 2012 | Absolute values |
| Total Cost | Total cost, including wages, intermediate goods, other expenses, excluding investments, in Swiss Francs. Reported balance sheet values for 2010 and 2011, projected values for 2012 | Absolute values |
| Purchase Price Response | Purchase price response to a 30% (1 St. Dev.) oil price shock: Average effect on all purchases of goods and services | $\leq -7.5\%$ to $\geq +7.5\%$ |
| Total Cost Response | Total cost response to a 30% (1 St. Dev.) oil price shock: Total cost response including wages, intermediate goods, other expenses, excluding investments | $\leq -7.5\%$ to $\geq +7.5\%$ |
| Domestic Sales Price Response | Domestic sales price response to a 30% (1 St. Dev.) oil price shock: Sales price response for sales in Switzerland | $\leq -7.5\%$ to $\geq +7.5\%$ |
| Foreign Sales Price Response | Foreign sales price response to a 30% (1 St. Dev.) oil price shock: Sales price response for sales outside Switzerland | $\leq -7.5\%$ to $\geq +7.5\%$ |
| Nominal Turnover Response | Nominal turnover response to a 30% (1 St.Dev.) oil price shock | $\leq -7.5\%$ to $\geq +7.5\%$ |
| Real Turnover Response | Nominal turnover divided by the weighted mean of domestic and foreign sales price responses. Weights are derived from Euro area and world export share variables. | |
| Value Added | Gross value added figures for 2011 for all Swiss industries based on international NACE classification scheme. Source: Swiss Statistical Office, www.bfs.admin.ch | Absolute values |

Responses based on judgement in June/July 2012 for expected effects by end of 2012 (6 months ahead) and by the end of 2013 (18 months) ahead. All responses have been transformed to continuous scale.

Aggregation scheme

We aggregate firms' responses by calculating a weighted mean,

$$\bar{y} = \sum_{i=1}^N w_i y_i.$$

where y_i is the response of firm $i = 1, \dots, N$ and w_i is its specific weight. The corresponding weighted standard deviation is then

$$\sigma_{y_i} = \sqrt{\sum_{i=1}^N (y_i - \bar{y})^2 w_i}$$

The weights w_i are derived from an aggregation scheme so that any coverage error does not induce a bias in the results and representativeness is ensured. Specifically,

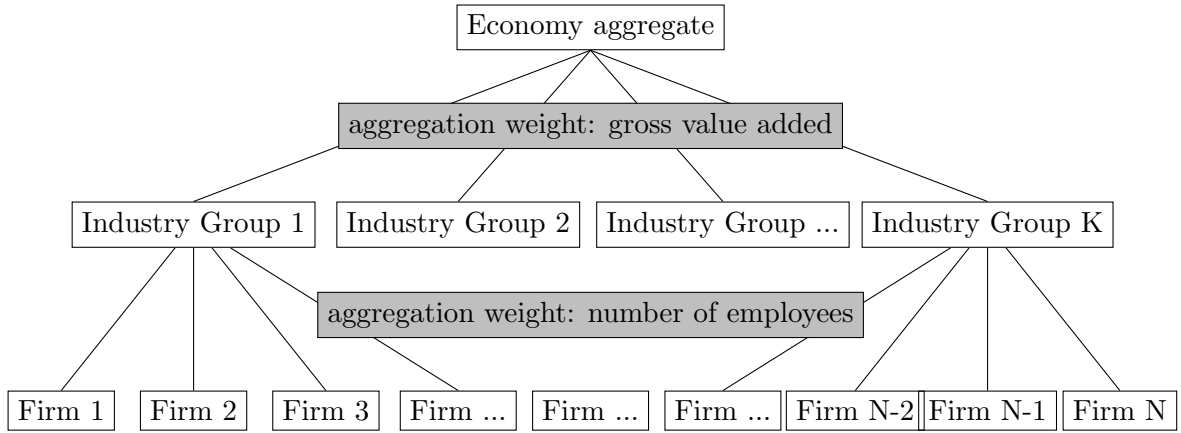
$$w_i = w_i^{Emp} * w_i^{VA},$$

where w_i^{Emp} is the number of employees of firm i divided by the cumulative number of employees of all firms within firm i 's industry group and where w_i^{VA} is the gross value added of firm i 's industry group divided by the cumulative gross value added of all industry groups together. The value added data have been taken from the 2011 Value Added Statistics of the Swiss Statistical Office. Figure 6.10 depicts the aggregation scheme. The cumulative gross value added of all industry groups has been adjusted by omitting those industry groups, for which we did not observe a sufficient number of firms employees. The aim was to have at least 30 observations within each group (with the exception of motor vehicles, furniture, and repair & installation works).

VAR companion model

To identify the effects of oil price shocks on the Swiss economy, we follow the strategy suggested by Kilian (2009), who used this approach to identify the effects of oil price

Figure 6.10: Aggregation scheme



shocks on the U.S. economy. In a first step, we identify oil supply, oil demand, and precautionary oil demand shocks at the international level using a VAR. To find the effects of these international oil price shocks on the Swiss economy, we use a distributed lag model of Swiss aggregated variables and the identified oil price shocks. We employ Bayesian methods to estimate the VAR and the distributed lag model. The Bayesian approach has the advantage that both models can be estimated in one step, which simplifies the assessment of uncertainty around our estimates.

World-wide model

Consider the following VAR:

$$y_t = c + \sum_{i=1}^{24} B_i y_{t-i} + u_t \quad (6.3)$$

with $y_t = [\Delta oilprod_t, react_t, oilprice_t]$, where $\Delta oilprod_t$ is world crude oil production, $react_t$ is the world real economic activity index provided by Kilian (2009)¹⁵, and $oilprice_t$ is the crude oil price (Brent). All variables are on a monthly basis ranging from April 1990 to April 2013. The reduced-form error u_t is normally distributed

¹⁵We use the updated version of the series which can be downloaded from Kilian's webpage.

with mean zero and variance covariance matrix Σ . The reduced-form error term can also be written as $u_t = A\epsilon_t$, where A is a lower triangular matrix obtained from a Cholesky decomposition of Σ . As in Kilian (2009) the vector of structural shocks ϵ_t consists of an oil supply shock, an aggregate demand shock, and an oil-specific demand shock.

Model for Switzerland

To calculate the effects of the oil price shocks identified in the VAR on Swiss GDP and inflation, consider the following regressions:

$$x_t = \alpha_j + \sum_{i=1}^{12} \beta_{ij} \xi_{jt-i} + e_{jt,x} \quad (6.4)$$

with

$$e_{jt,x} = \sum_{i=1}^3 \phi_{ij,x} e_{jt-i,x} + \nu_{jt,x}, \quad (6.5)$$

where x_t is turnover, $\nu_{jt,x}$ is a normally distributed disturbance with mean zero and standard deviation $\sigma_{j,x}$, and ξ_t is the vector of structural shocks at quarterly frequency.¹⁶ Because Swiss GDP is available only at a quarterly frequency the oil shocks are aggregated from a monthly to a quarterly frequency using the following equation:

$$\xi_t = \frac{1}{3} \sum_{m=1}^3 \epsilon_{mt}, \quad (6.6)$$

where m denotes months at each quarter t . The regressions for the effect of the oil price shocks on Swiss PPI inflation can be expressed as

$$\pi_t = \delta_j + \sum_{i=1}^{12} \theta_{ji} \xi_{jt-i} + e_{jt,\pi}, \quad (6.7)$$

¹⁶Equations (6.5) and (6.8) deviates from the set-up in Kilian (2009). Kilian (2009) does not model serial correlation in the error terms.

with

$$e_{jt,\pi} = \sum_{i=1}^3 \phi_{ij,\pi} e_{jt-i,\pi} + \nu_{jt,\pi}, \quad (6.8)$$

where π_t denotes Swiss PPI inflation and $\nu_{jt,\pi}$ is a normally distributed disturbance with mean zero and standard deviation $\sigma_{j,\pi}$. As discussed in Section 4.3, we repeat the analysis using Swiss manufacturing real turnover instead of Swiss GDP.

Estimation

The model is estimated using Bayesian methods, more precisely the Gibbs sampling procedure. The Gibbs sampler consists of the following blocks. In a first step, conditional on the variance covariance matrix Σ draw the VAR coefficients B_1, \dots, B_{24}, c from a multivariate normal distribution. Conditional on the VAR coefficients, obtain draws for Σ from an inverted Wishart distribution (see, e.g., Karlsson and Kadiyala, 1997). In the next step, draw the $\beta_{1j}, \dots, \beta_{12j}, \alpha_j$ (or $\theta_{1j}, \dots, \theta_{12j}, \delta_j$) conditional on the VAR parameters and the particular parameters of the serially correlated error equations from a multivariate normal distribution. Conditional on the coefficients of the distributed lag model and the variance of the disturbance equations, draw the coefficients $\phi_{1j,x}, \dots, \phi_{3j,x}$ (or $\phi_{1j,\pi}, \dots, \phi_{3j,\pi}$) from a multivariate normal distribution. Finally, conditional on the coefficients of the distributed lag model and the serially correlated error equations draw the variance $\sigma_{j,x}$ ($\sigma_{j,\pi}$) from an inverted gamma distribution (see, e.g., Chib, 1993 for the last three blocks). After the estimation of the last block, we start the next iteration step from the first block again by conditioning on the last iteration step. These iterations have the Markov property: as the number of steps increases, the conditional posterior distributions of the parameters and the factor converge to their marginal posterior distributions at an exponential rate (see Geman and Geman, 1984).

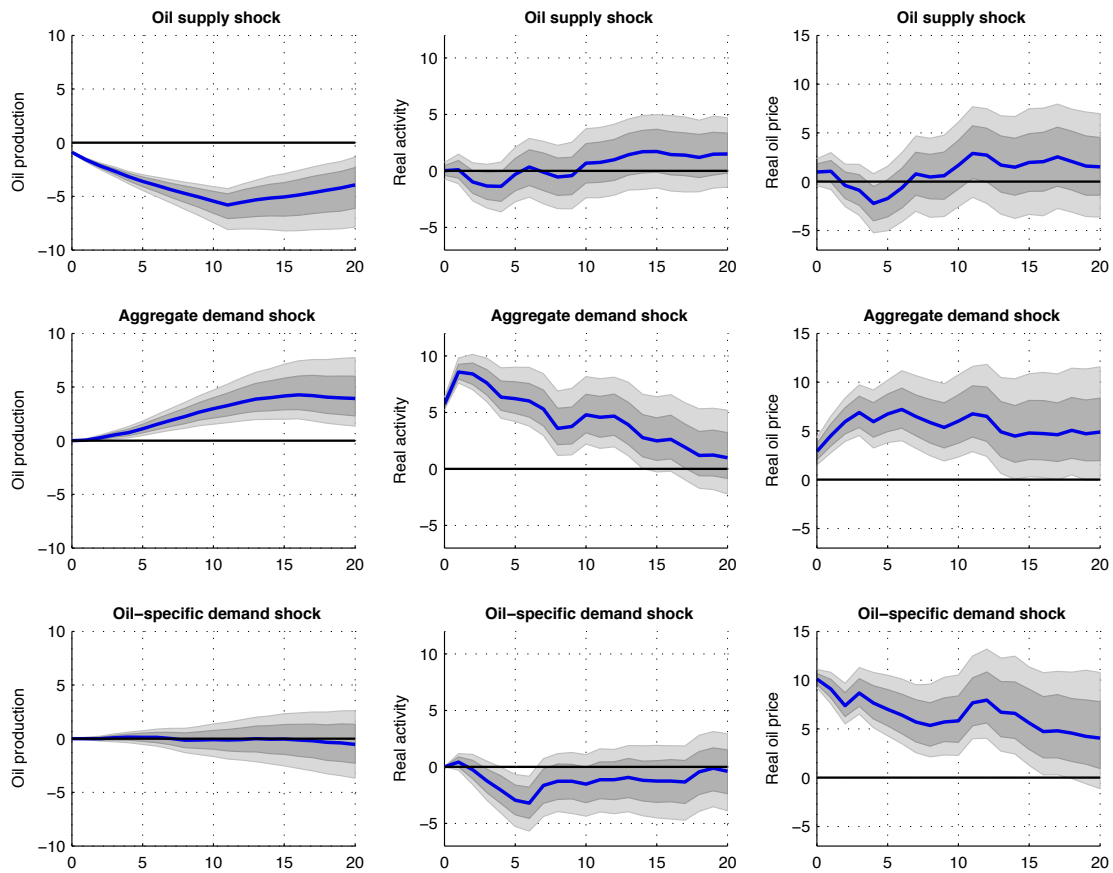
The priors specified for the model's parameters are all extremely diffuse. The prior for the VAR parameters follow a Normal-Wishart distribution and the one for the ADL parameters follow a Normal-Gamma distribution.

The results were computed using 500.000 draws from the Gibbs sampler. The first 400.000 draws were discarded as burn-in. From the remaining 100.000 draws we

have saved each 100th draw resulting into a sample of 1,000 draws. Convergence of the Gibbs sampler was checked by inspecting recursive mean plots of the parameters and by starting from different initial values and by comparing the results.


Figure 6.11 shows impulse responses in the world-wide model. Figure 6.12 displays the effects of oil price shocks on Swiss GDP and PPI while Figure 6.13 presents the results for Swiss manufacturing real turnover and PPI inflation.

Figure 6.11: Impulse response world-wide



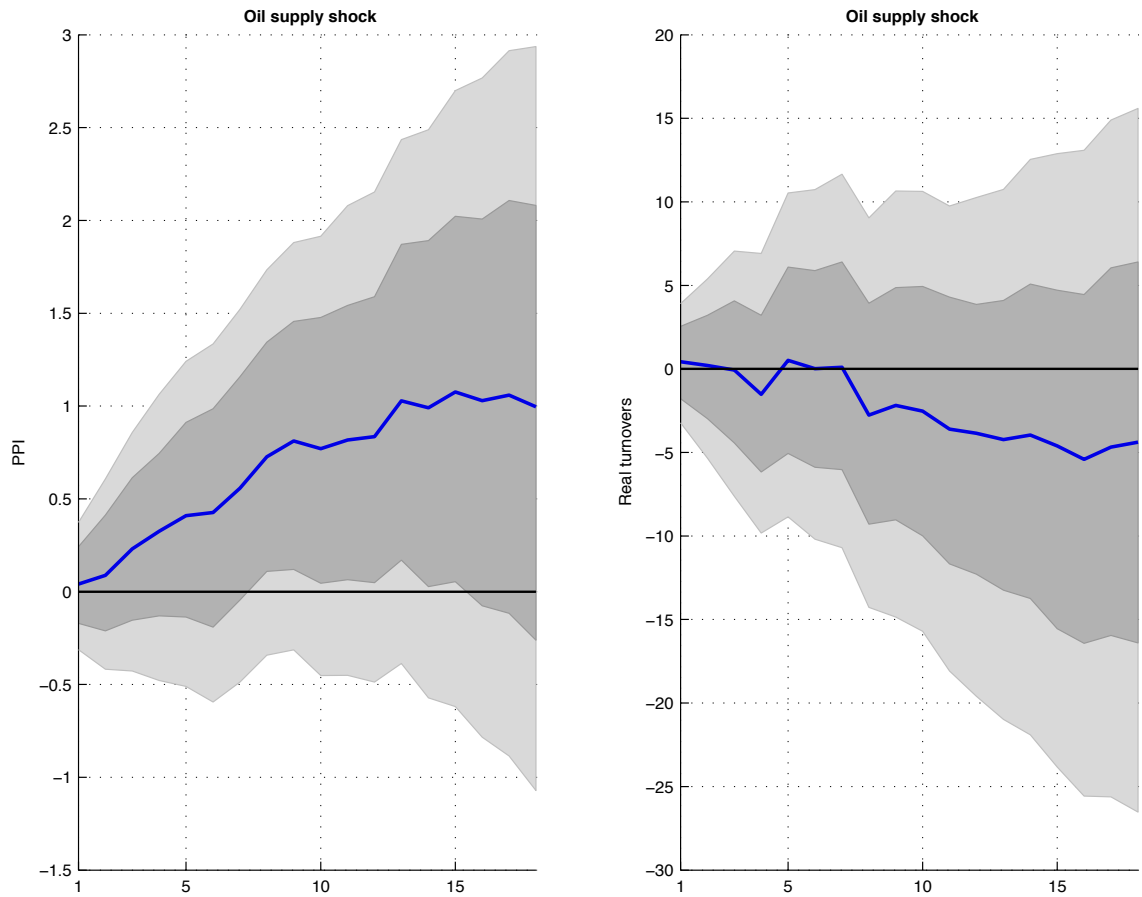
Impulse responses for a one-standard deviation shock in the world wide model. The dark gray area indicates the 16th and 84th percentiles of the impulse responses and the light gray area indicates the 5th and 95th percentiles.

Figure 6.12: Impulse response Swiss level



Responses of Swiss GDP and PPI level to each structural shock. The dark gray area indicates the 16th and 84th percentiles of the impulse responses and the light gray area indicates the 5th and 95th percentiles.

Figure 6.13: Responses of Swiss manufacturing real turnover



Responses of Swiss manufacturing real turnover and PPI level to each structural shock. The dark gray area indicates the 16th and 84th percentiles of the impulse responses and the light gray area indicates the 5th and 95th percentiles.

Appendix C

Appendix to Chapter 4

Validity of the Survey Experiment

In a previous section, we have addressed various aspects of possible survey errors described in the handbooks of Dillman et al. (2014). An elaborate discussion of how the results of our scenario survey are potentially affected by sampling error, measurement error due to interviewers or respondents, coverage error, non-response and post-survey error can also be found in the appendix to Drechsel et al. (2015b). This section focuses on potential selection issues caused by non-response if non-response is dependent on the specific question and context. Hence, the following paragraphs present a specific discussion of our tests for non-response bias.

Similar to self-selection issues, our results could potentially be biased by systematic non-response. If our additional questionnaire was highly relevant to a particular group of firms, these firms are very likely to respond, while firms for whom the questionnaire is less relevant systematically choose to drop out. If our questionnaire was only relevant to firms with large export shares, our results might over-estimate the effect of an exchange rate shock on the entire economy, as highly export-dependent firms are over-represented in the remaining sample.

To strengthen our trust in the fact that the estimation results are not biased by systematic non-response, we apply the “surrogate” approach of Wallace and Mellor (1988). We compare the firms that responded on time (i.e., by July 9, 2012) with firms that did not respond before the official due date. After the due date, we called companies who had not responded and urged them to participate.

Hence, we can interpret respondents who answered after the official deadline as a sample of the non-response population. Following Wallace and Mellor (1988), we create two subsamples by selecting the first 50 observations and the last 50 observations. Given that participants answered in random order, both subsamples should be random draws from the total population and thus should not differ in their means and distribution. Based on this idea we test for differences in export share of both subsamples. Admittedly, this could be done for differences in any of our variables; however, we focus on the export share because of its focal role. We also perform a Kolmogorov-Smirnov test to check whether distributions of both subsamples are equal. Table 6.5 shows the results of a t-test for equality of export share

means among late and early respondents. The difference in mean export shares are not significant at the 10 percent level.

Table 6.5: Mean export shares of early and late respondents

| Test | Mean export share of early respondents | Mean export share of late respondents | p-value |
|--------|---|--|---------|
| t-test | 15.10638 | 22.32558 | 0.1866 |

The sample mean provides a first glance at the export share variable but does not fully represent differences in sampling distributions. Thus, we also look at the sampling distributions of both samples. The Kolmogorov-Smirnov (KS) test fails to reject the hypothesis of equal distributions for both the weighted and the unweighted sample. Table 6.6 summarizes the results of the KS test. Note that the p-value is not exact if ties occur in the KS-test. However, with the p-value not close to alpha this does not seem to be an issue here. Results seem not to be biased by the varying relevance of questions to firms.

Table 6.6: Kolmogorov-Smirnov test results for equality of distributions

| Test | Test statistic D | p-value |
|---------|------------------|---------|
| KS-test | 0.1979 | 0.3426 |

Enterprise panel

The enterprise panel of KOF Swiss Economic Institute is based on a sample of 7000 firms taken from the Business Register (BR) of the Federal Statistical Office. The sample, which covers the manufacturing, construction, and the commercial area of the service sector, is stratified according to sectors and sector-specific variables and is adjusted regularly. The corresponding address database, which in addition to numerous structural features of the firms, contains all information that is required to ensure that panel surveys run smoothly (specific contact points, checking of replies, recalls, and incomplete surveys etc.), was developed continuously and updated regularly throughout the project.

Based on the KOF enterprise panel, regular collections of data on the structural changes, innovative activities and investment plans of Swiss companies are conducted. The enterprise panel draws from a population of 60270 companies. These companies are part of the entire collection of firms within Switzerland. Out of all firms within Switzerland, only those with a business register number (assigned in 2001) and with at least 5 employees have been selected. Companies belonging to agricultural activities or public administration have been included. In contrast to the sample used for structural investigations of the Swiss economy, the utilized investment sample also contains firms active in education, health, waste disposal, entertainment, cultural and sports activities.

Questionnaire

The questionnaire used for this paper was attached to KOF Investment Survey and was developed in the Macrolab project. The entire questionnaire can be found in the appendix of chapter 3 which makes use of the same survey.

Data

Are detailed description of the variables obtained from the survey experiment and external sources can be found in the appendix to Chapter 3. The aggregation from the firm level to industry and sector levels follows the same aggregation scheme used

in Chapter 3. Further, trade surplus is calculated as the ratio between the difference of nominal exports ($X_{i_{s=6}}^{Euro,world}$) and nominal imports ($IM_{i_{s=6}}^{Euro,world}$) and nominal unconditional turnover ($Y_{i_{s=6}}^{unc}$) of firm i at horizon $s = 6$ months:

$$\frac{X_{i_{s=6}}^{Euro,world} - IM_{i_{s=6}}^{Euro,world}}{Y_{i_{s=6}}^{unc}}$$

with firms $i = 1, \dots, I$. Assuming that the export share stays constant nominal exports $X_{i_{s=6}}^{Euro,world}$ are calculated as

$$X_{i_{s=6}}^{Euro,world} = \frac{X_{i_{s=0}}^{Euro,world}}{Y_{i_{s=0}}} \times Y_{i_{s=6}}^{unc}$$

$\frac{X_{i_{s=0}}^{Euro,world}}{Y_{i_{s=0}}}$ is the mid-2012 (Euro or world) export share with respect to turnover and has been provided by the participating firms. In the same manner, nominal imports $IM_{i_{s=6}}^{Euro,world}$ are approximated as

$$IM_{i_{s=6}}^{Euro,world} = \frac{IM_{i_{s=0}}^{Euro,world}}{C_{i_{s=0}}} \times C_{i_{s=6}}^{unc}$$

where $C_{i_{s=6}}^{unc}$ are the unconditional costs forecast of firm i for 2012. $\frac{IM_{i_{s=0}}^{Euro,world}}{C_{i_{s=0}}}$ is the mid-2012 import share with respect to total costs and has been provided by the participating firms.

Appendix D

Appendix to Chapter 6

Installation notes

R stable version

The stable version of the `timeseriesdb` R package itself can be downloaded and installed from CRAN (R's official repository). The package source as well as binaries for Windows and OS X are available from CRAN. The package can be installed following R's standard procedure to install packages by either running:

```
install.packages("timeseriesdb")
```

or using R's GUI.

R developer version

The developer version of `timeseriesdb` can be obtained from github.com/mbannert/timeseriesdb. The most convenient way to install the latest developer version from inside an R session is to use the `devtools` package (Wickham and Chang, 2014):

```
library(devtools)
install_github('mbannert/timeseriesb')
```

PostgreSQL

Because `timeseriesdb` depends on `RPostgreSQL` to connect to PostgreSQL databases, the user needs to make sure that the PostgreSQL's own library and header files are present and can be found by `RPostgreSQL`. For Windows, this library called `libpq` and is attached to the `RPostgreSQL` package and will thus be installed with the `RPostgreSQL` package. Hence, Windows users should not experience any difficulties.

For OS X and Linux, the installation is a bit different when `libpq` is not present. For some Linux distributions, the corresponding library can be obtained with the `postgresql-devel` package. Similarly, on OS X, the user needs to make sure that `libpq` is present and can be found by `Rpostgresql`. It is recommend to use the

homebrew package manager running ‘brew install postgresql’. OS X and Linux users should note that previously installed versions of PostgreSQL may not contain the libraries provided by `postgresql-devel` package.

Database setup

If you do not have a PostgreSQL database that contains a timeseries schema that suits `timeseriesdb`, create a schema called `timeseries` and run `setup.sql`. The file is located in `inst/sql` of your package folder. Start a psql client console from the `inst/sql` directory and run: `\i setup.sql`. If you are not familiar with running a PostgreSQL console, copy and paste the content of that file to the SQL window of your favorite GUI tool (e.g., PGAdmin) and run it.

Benchmarks

Because `timeseriesdb` aims at storing a large number of time series, it is likely that `timeseriesdb` is used in bulk processes, such as reading thousands of time series into an R session, process them in R, and write results back to the database.¹⁷ When processing a larger amount of time series, the speed of reading and writing to the database can become a substantial part of a process’ total runtime. Thus `timeseriesdb` strives to speed up reading and writing. The benchmark example shown in the table below compares three different ways to read from the database.

Table 6.7: Benchmark: Reading 100 time series, 100 replications

| | test | replications | elapsed | relative |
|---|---|--------------|---------|----------|
| 1 | <code>lapply(paste0("ts", 1:100), loopRead, con = con)</code> | 100 | 79.38 | 15.41 |
| 2 | <code>readAndSplit(paste0("ts", 1:100), con)</code> | 100 | 21.51 | 4.18 |
| 3 | <code>readTimeSeries(paste0("ts", 1:100), con)</code> | 100 | 5.15 | 1.00 |

The first line shows benchmark results for a simple read function that reads a single time series from the database given its key. Using a function that reads a

¹⁷Storing and reading even larger chunks of time series (i.e., millions of time series with hundreds of observations) can cause memory leaks. The latest developer version of `timeseriesdb` can also handle such larger amounts of data by splitting the data according to C stack size.

single series and loops over it would be the most intuitive solution for most users. However, doing so creates a new `SELECT` statement for every query executed. Plus, a new database transaction is started for every iteration. Particularly using multiple `SELECT` statements slows reading down considerably. Hence, the approach shown in line 2 of the benchmark resolves the *hstore* data type and returns all observations from every time series to an R *data.frame*. This *data.frame* is then split by key and sorted by time using the `data.table` package (Dowle et al., 2014). The `data.table` package moves the split operation to C++ and is therefore able to speed up the read process substantially; however, splitting the data by key is still computationally costly. The third line displayed in the benchmark is the version that is currently used within the `timeseriesdb` package. The function `readTimeSeries` makes use of a view that exposes entire time series records including key and frequency in a *JSON* object. Therefore, the current read function only needs one `SELECT` statement and uses a `WHERE IN` clause. With the help of `RJSONIO`, the *JSON* objects can be resolved easily without using costly split actions. That makes the current version more than 15 times faster than the basic solution and still more than four times faster than splitting the set at the R, respectively, the C++ level.

Bibliography

- Abberger, K., M. Bannert, and A. Dibiasi (2014). Metaumfrage im dienstleistungssektor. *KOF Analysen* 8(2), 51–62.
- Abberger, K., M. Birnbrich, and C. Seiler (2011). Der test des tests im handel—eine metaumfrage zum ifo konjunkturtest. *ifo Schnelldienst* 62(21), 34–41.
- Abrahamsen, Y. and B. Simmons-Süer (2011, March). Die Wechselkursabhängigkeit der Schweizer Wirtschaft. *KOF Analysen* 5(1), 61–74.
- Allaire, J., J. McPherson, Y. Xie, H. Wickham, J. Cheng, and J. Allen (2014). *rmarkdown: Dynamic Documents for R*. R package version 0.5.3.1.
- Angrist, J. D. and J.-S. Pischke (2010). The Credibility Revolution in Empirical Economics: How Better Research Design is Taking the Con out of Econometrics. *Journal of Economic Perspectives* 24(2), 3–30.
- Bannert, M. (2015). *timeseriesdb: Store and Manage Time Series in a Database*. R package version 0.1.
- Barabas, J. and J. Jerit (2010, May). Are Survey Experiments Externally Valid? *American Political Science Review* 104(02), 226–242.
- Beck, K., U. Kunze, and W. Oggier (2004). Selbstdispensation: Kostentreibender oder Kostendämpfender Faktor? *Managed Care*, 33–36.
- Bengtsson, H. (2015). *R.matlab: Read and Write MAT Files and Call MATLAB from Within R*. R package version 3.2.0.

- Bils, M. and P. J. Klenow (2004). Some Evidence on the Importance of Sticky Prices. *Journal of Political Economy* 112(5), 947–985.
- Bivand, R. S., E. Pebesma, and V. Gomez-Rubio (2013). *Applied spatial data analysis with R, Second edition*. Springer, NY.
- Branchi, M., H. C. Dieden, W. Haine, C. Horvath, A. Kanutin, and L. Kezber (2007). Quarterly gdp revisions in g-20 countries: Evidence from the 2008 financial crisis. *ECB Occasional Paper* (74).
- Buckheit, J. B. and D. L. Donoho (1995). Wavelab and reproducible research. In *Wavelets in Statistics, Lecture Notes in Statistics*, pp. pp. 55–82. Springer.
- Busato, A., P. Matter, B. Künzi, and D. C. Goodman (2010). Supply sensitive services in swiss ambulatory care: An analysis of basic health insurance records for 2003-2007. *BMC Health Services Research* 10.
- Cameron, C. A. and P. K. Trivedi (2013). Cambridge University Press, New York.
- Campa, J. M. and L. S. Goldberg (2005). Exchange Rate Pass-Through into Import Prices. *The Review of Economics and Statistics* 87(4), 679–690.
- Chang, W., J. Cheng, J. Allaire, Y. Xie, and J. McPherson (2015). *shiny: Web Application Framework for R*. R package version 0.11.1.
- Charness, G., U. Gneezy, and M. A. Kuhn (2012). Experimental Methods: Between-Subject and Within-Subject Design. *Journal of Economic Behavior & Organization* 81(1), 1–8.
- Chib, S. (1993). Bayes regression with autoregressive errors : A Gibbs sampling approach. *Journal of Econometrics* 58(3), 275–294.
- Chou, Y., W. C. Yip, C.-H. Lee, N. Huang, Y.-P. Sun, and H.-J. Chang (2003). Impact of separating drug prescribing and dispensing on provider behaviour: Taiwan’s experience. *Health Policy and Planning* 18(3), 316–329.
- Cochrane, W. G. (1977). *Sampling Techniques*. John Wiley & Sons, Hoboken.

- Conway, J., D. Eddelbuettel, T. Nishiyama, S. K. Prayaga, and N. Tiffin (2013). *RPostgreSQL: R interface to the PostgreSQL database system*. R package version 0.4.
- Couper, M.P., Singer, E., Conrad, F.G., and Groves, R.M (2008). Risk of disclosure, perceptions of risk, and concerns about privacy and confidentiality as factors in survey participation. *Journal of Official Statistics* 24, 255–275.
- Cox, D. and W. Smith (1989). An outline of generalized linear models. In P. McCullagh and J. Nelder (Eds.), *Generalized Linear Models*, pp. 21–47. Chapman and Hall, New York.
- Dalenius, T. and M. Gurney (1951, January). The Problem of Optimum Stratification. II. *Scandinavian Actuarial Journal* 1951(1-2), 133–148.
- Denis Mukhin, D. A. J. and J. Luciani (2014). *ROracle: OCI based Oracle database interface for R*. R package version 1.1-12.
- Dillman, D. A. (2000). *Mail and internet surveys: The tailored design method*, Volume 2. Wiley New York.
- Dillman, D. A., J. D. Smyth, and L. M. Christian (2014). *Internet, Phone, Mail and Mixed Mode Surveys - The Tailored Design Method* (4th ed.). John Wiley & Sons, Hoboken, New Jersey.
- Dowle, M., T. Short, S. Lianoglou, A. S. with contributions from R Saporta, and E. Antonyan (2014). *data.table: Extension of data.frame*. R package version 1.9.4.
- Drechsel, D., H. Mikosch, S. Sarferaz, and M. Bannert (2015a). How are Firms Affected by Exchange Rate Shocks? Evidence from Survey Based Impulse Responses. *KOF Working Paper Series* (371).
- Drechsel, D., H. Mikosch, S. Sarferaz, and M. Bannert (2015b). *Survey based Impulse Responses*. mimeo, ETH Zurich.

- Druckman, J. N., D. P. Green, J. H. Kuklinski, and A. Lupia (2006). The Growth and Development of Experimental Research in Political Science. *American Political Science Review* 100(4), 627–635.
- Emery, J., D. Lima, J. Lewis, and B. Sunderland (2009). A systematic review of the literature comparing the practices of dispensing and non-dispensing doctors. *Health Policy* 92, 1–9.
- European Commission (2007). The Joint Harmonised EU Programme of Business and Consumer Surveys. Technical report, European Commission, Directorate-General for Economic and Financial Affairs.
- Faruqee, H. (2004, January). Exchange Rate Pass-Through in the Euro Area; The Role of Asymmetric Pricing Behavior. IMF Working Papers 04/14, International Monetary Fund.
- Filippini, M., F. Heimsch, and G. Masiero (2013, March). Antibiotic consumption and the role of dispensing physicians. *CEPRA working paper* (1302).
- Friedman, J., T. Hastie, and R. Tibshirani (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 1–22.
- Fukunaga, I., N. Hirakata, and N. Sudo (2010). The Effects of Oil Price Changes on the Industry-Level Production and Prices in the U.S. and Japan. *NBER Working Paper* 15791, 1–52.
- Gaines, B. J., J. H. Kuklinski, and P. J. Quirk (2006). The Logic of the Survey Experiment Reexamined. *Political Analysis* 15(1), 1–20.
- Gandrud, C. (2013). *Reproducible Research with R and R Studio*. Chapman and Hall/CRC.
- Geman, S. and D. Geman (1984). Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6(6), 721–41.

- Gilbert, P. (2013a). *TSdbi: TSdbi (Time Series Database Interface)*. R package version 2013.9-1.
- Gilbert, P. (2013b). *TSMYSQL: TSdbi extensions for MySQL*. R package version 2013.9-1.
- Gilbert, P. (2013c). *TSPostgreSQL: TSdbi extensions for PostgreSQL*. R package version 2013.9-1.
- Gilens, M. (2002). An Anatomy of Survey-Based Experiments. In J. Manza, F. L. Cook, and B. J. Page (Eds.), *Navigating Public Opinion: Polls, Policy, and the Future of American Democracy*, pp. 232–250. Oxford University Press, Oxford.
- Gopinath, G., O. Itskhoki, and R. Rigobon (2010). Currency Choice and Exchange Rate Pass-Through. *American Economic Review* 100(1), 304–36.
- Gopinath, G. and R. Rigobon (2008). Sticky Borders. *The Quarterly Journal of Economics* 123(2), 531–575.
- Greer, S. L. and P. D. Jacobson (2010). Health care reform and federalism. *Journal of Health Politics, Policy and Law* 35(2), 203–226.
- Groves, R. M. (1989). *Survey Errors and Survey Costs*. John Wiley & Sons, Hoboken.
- Groves, R. M. and M. P. Couper (1998). *Nonresponse in household interview surveys*. John Wiley & Sons.
- Groves, R. M., D. Dillman, J. L. Eltinge, and R. J. Little (2002). *Survey nonresponse*. Wiley New York.
- Guterbock, T. M. and S. L. Nock (2010). Survey Experiments. In J. Wright and P. Marsden (Eds.), *Handbook of Survey Research*, pp. 837–864. Emerald Group Publishing, Bingley.
- Hahn, E. (2003, July). Pass-through of external shocks to euro area inflation. Working Paper Series 0243, European Central Bank.

- Hahn, E. (2007). The Impact Exchange Rate Shocks on Sectoral Activity and Prices in the Euro Area. Working Paper Series 0243, European Central Bank.
- Hamilton, J. D. (1988). A Neoclassical Model of Unemployment and the Business Cycle. *Journal of Political Economy* 96(3), 593–617.
- Hamilton, J. D. (1994). *Time Series Analysis*. Princeton University Press, Princeton.
- Hamilton, J. D. (2008). Understanding Crude Oil Prices. *The Energy Journal* 30(2), 179–206.
- Harris-Kojetin, B. and C. Tucker (1999). Exploring the relation of economic and political conditions with refusal rates to a government survey. *JOURNAL OF OFFICIAL STATISTICS-STOCKHOLM-* 15, 167–184.
- Herrera, A. M. (2008). Oil Price Shocks , Inventories and Macroeconomic Dynamics. *Michigan State University Working Paper* , 1–64.
- Horvitz, D. G. and D. J. Thompson (1952). A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association* 47(260), 663–685.
- Iizuka, T. (2007). Experts’ agency problems: evidence from the prescription drug market in japan. *The RAND Journal of Economics* 38(3), 844–862.
- Jiménez-Rodríguez, R. (2011, March). Macroeconomic Structure and Oil Price Shocks at the Industrial Level. *International Economic Journal* 25(1), 173–189.
- Kaiser, B. and C. Schmid (2014). Does physician dispensing increase drug expenditures? empirical evidence from switzerland. *Health Economics*.
- Karlsson, S. and K. R. Kadiyala (1997). Numerical Methods for Estimation and Inference in Bayesian Var-Models. *Journal of Applied Econometrics* 12(12), 99–132.
- Kilian, L. (2008, November). The Economic Effects of Energy Price Shocks. *Journal of Economic Literature* 46(4), 871–909.

- Kilian, L. (2009). Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market. *American Economic Review* 99(3), 1053–1069.
- Kinder, D. R. (2007, March). Curmudgeonly Advice. *Journal of Communication* 57(1), 155–162.
- Koenker, R. and A. Zeileis (2009). On reproducible econometric research. *Journal of Applied Econometrics* 24(5), 833–847.
- Kuziemko, I., M. I. Norton, E. Saez, and S. Stantcheva (2015). How elastic are preferences for redistribution? evidence from randomized survey experiments. *American Economic Review* 105(4), 1478–1508.
- Lagalo, L. G. (2011). Separating Demand and Supply Shocks in the Oil Market – An Analysis Using Disaggregated Data. *University of Michigan Working Paper* , 1–34.
- Lang, D. T. (2014). *RJSONIO: Serialize R objects to JSON, JavaScript Object Notation*. R package version 1.3-0.
- Lassmann, A. (2013). Exchange Rate Transmission and Export Activity at the Firm Level. KOF Working Paper Series 331, ETH KOF.
- Leamer, E. E. (2010). Tantalus on the Road to Asymptopia. *Journal of Economic Perspectives* 24(2), 31–46.
- Lee, K. and S. Ni (2002). On the Dynamic Effects of Oil Price Shocks: A study Using Industry Level Data. *Journal of Monetary Economics* 49(4), 823–852.
- Leeper, T. J. (2014, June). Archiving reproducible research with r and dataverse. *The R Journal* 6(1), 151–158.
- Little, R. J. and D. B. Rubin (2002). *Statistical Analysis with missing data*.

- Liu, Y.-M., Y.-H. K. Yang, and C.-R. Hsieh (2009). Financial incentives and physiciansâ€™TM prescription decisions on the choice between brand-name and generic drugs: Evidence from taiwan. *Journal of Health Economics* 28(2), 341 – 349.
- Lundin, D. (2000). Moral hazard in physician prescription behavior. *Journal of Health Economics* 19(5), 639 – 662.
- Luxen, D. and C. Vetter (2011). Real-time routing with openstreetmap data. In *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, GIS '11, New York, NY, USA, pp. 513–516. ACM.
- McCullough, B. D. (2009). Open access economics journals and the market for reproducible economic research. In *Economic Analysis and Policy* 39.1, pp. pp. 117–126.
- McCullough, B. D. and H. D. Vinod (2003). Verifying the solution from a nonlinear solver: A case study. *American Economic Review* 93(3), 873–892.
- Meeus, J. (1999). *Astronomical Algorithms*. Willmann-Pell.
- Nakamura, E. and J. Steinsson (2008). Five Facts About Prices: A Reevaluation of Menu Cost Models. *Quarterly Journal of Economics* 123(4), 1415–1464.
- NIMA (2000). Department of defense world geodetic system 1984. Technical report, National Imagery and Mapping Agency.
- Ooms, J. (2014a). The jsonlite package: A practical and consistent mapping between json data and r objects. *arXiv:1403.2805 [stat.CO]*.
- Ooms, J. (2014b). *opencpu: The OpenCPU system for embedded scientific computing and reproducible research*. R package version 1.4.5.
- Pebesma, E. and R. Bivand (2005). Classes and methods for spatial data in r. *R News* 2(5).

- Peersman, G. and I. Van Robays (2012, September). Cross-country Differences in the Effects of Oil Shocks. *Energy Economics* 34(5), 1532–1547.
- pharmaSuisse (2014). Swiss pharmacists' association annual report 2013. Technical report, pharmaSuisse.
- Porst, R. and C. von Biel (1995). Wären sie vielleicht bereit, sich gegebenenfalls noch einmal befragen zu lassen?: oder: Gründe für die teilnahme an panelbefragungen. *ZUMA Arbeitsbericht*.
- R Core Team (2015). *foreign: Read Data Stored by Minitab, S, SAS, SPSS, Stata, Systat, Weka, dBase, ...* R package version 0.8-63.
- Reich, O., C. Weins, C. Schusterschitz, and M. Thöni (2012). Exploring the disparities of regional health care expenditures in switzerland: some empirical evidence. *The European Journal of Health Economics* 13(2), 193–202.
- Reips, U.-D. (2009). Schöne neue forschungswelt: Zukunftstrends [beautiful new world of research: Future trends]. nicht-reaktive erhebungsverfahren. In *GESIS Schriftenreihe*, Volume 1, pp. 129–138.
- Rischatsch, M. and M. Trottmann (2009). Physician dispensing and the choice between generic and brand-name drugs – Do margins affect choice? *SOI - Working Papers* (0911).
- Rudebusch, G. D. (1998). Do Monetary Policy in a VAR Make Sense ? *International Economic Review* 39(4), 907–931.
- Ryan, J. A. and J. M. Ulrich (2013). *xts: eXtensible Time Series*. R package version 0.9-7.
- Sax, C. (2014). *seasonal: R interface to X-13ARIMA-SEATS*. R package version 0.62.1.
- Seiler, C. (2012). On the robustness of the balance statistics with respect to nonresponse. Technical report, ifo Working Paper.

- Seiler, C. (2013). *Nonresponse in business tendency surveys*. Ph. D. thesis, lmu.
- Shrestha, M. and M. Marini (2013). Quarterly gdp revisions in g-20 countries: Evidence from the 2008 financial crisis. *IMF Working Paper*.
- Sims, C. A. (2010). But Economics Is Not an Experimental Science. *Journal of Economic Perspectives* 24(2), 59–68.
- Sniderman, P. M. and D. B. Grob (1996, August). Innovations in Experimental Design in Attitude Surveys. *Annual Review of Sociology* 22(1), 377–399.
- Spitzmüller, C., D. M. Glenn, M. M. Sutton, C. D. Barr, and S. G. Rogelberg (2007). Survey nonrespondents as bad soldiers: Examining the relationship between organizational citizenship and survey response behavior. *International Journal of Selection and Assessment* 15(4), 449–459.
- Stieger, S. and U.-D. Reips (2010). What are participants doing while filling in an online questionnaire: A paradata collection tool and an empirical study. *Computers in Human Behavior* 26(6), 1488 – 1495. Online Interactivity: Role of Technology in Behavior Change.
- Stock, J. H. (2010). The Other Transformation in Econometric Practice: Robust Tools for Inference. *Journal of Economic Perspectives* 24(2), 83–94.
- Stokes, H. H. (2004). On the advantage of using two or more econometric software systems to solve the same problem. *Journal of Economic and Social Measurement* 29, 307–320.
- Tomaskovic-Devey, D., J. Leiter, and S. Thompson (1994). Organizational survey nonresponse. *Administrative Science Quarterly*, 439–457.
- Trap, B. (1997). Practices of dispensing doctors - drug use and health economics. *International Conferences on Improving Use of Medicines*.

- Trottmann, M., M. Früh, O. Reich, and H. Telser (2015). Auswirkungen der medikamentenabgabe durch die ärzteschaft (selbstdispensation) auf den arzneimittelkonsum und die kosten zu lasten der okp. Technical report, Polynomics, Studie im Auftrag des Bundesamtes für Gesundheit.
- Uhlmann, B. (2013). Federalism and pharmaceutical self-dispensing policy in switzerland: a comparative analysis of member states' preferences at national level, and vested interests influenced practices at cantonal level. *1st International Conference on Public Policy, 26-28 June – Grenoble*.
- Venables, W. N. and B. D. Ripley (2002). *Modern Applied Statistics with S* (Fourth ed.). New York: Springer. ISBN 0-387-95457-0.
- Vincenty, T. (1975). Direct and inverse solutions of geodesics on the ellipsoid with application of nested equations. (176), pp88–93.
- Wallace, R. S. O. and C. J. Mellor (1988). Nonresponse Bias in Mail Accounting Surveys: A Pedagogical Note. *British Accounting Review* 1988(20), 131–139.
- Weisberg, H. F. (2005). *The Total Survey Error Approach: A Guide To The New Science Of Survey Research*. Chicago: University of Chicago Press.
- Wickham, H. and W. Chang (2014). *devtools: Tools to make developing R code easier*. R package version 1.6.1.
- Wickham, H. and K. Müller (2014). *DBI: R Database Interface*. R package version 0.3.1.
- Willimack, D. K., E. Nichols, and S. Sudman (2002). Understanding unit and item nonresponse in business surveys. *Survey nonresponse*, 213–227.
- Xie, Y. (2013). *Dynamic Documents with R and knitr*. Chapman and Hall/CRC. ISBN 978-1482203530.
- Xie, Y. (2014). A comprehensive tool for reproducible research in r. In F. L. Victoria Stodden and R. D. Peng (Eds.), *Implementing Reproducible Computational Research*. Chapman and Hall/CRC. ISBN 978-1466561595.

Xie, Y. (2015). *knitr: A General-Purpose Package for Dynamic Report Generation in R*. R package version 1.9.

Yang, F., E. Tschetter, X. Léauté, N. Ray, G. Merlino, and D. Ganguli. *white paper*:

Zeileis, A. and G. Grothendieck (2005). zoo: S3 infrastructure for regular and irregular time series. *Journal of Statistical Software* 14(6), 1–27.

Zeileis, A., C. Kleiber, and S. Jackman (2008). Regression models for count data in r. *Journal of Statistical Software* 27(8), 1–25.

Imprint

Publisher

KOF Swiss Economic Institute, ETH Zurich

© 2016 Matthias Bannert

KOF Dissertation Series

- No 1 Causes and Consequences of Democracy and Terror**
Martin Gassebner
- No 2 The Impact of News on Expectation Formation in Monetary Policy, Business Cycles and Inflation Forecasts**
Michael J. Lamla
- No 3 Micro Price Adjustments and Macro Inflation Dynamics**
Sarah M. Lein
- No 4 Essays on Inflation Expectation Formation**
Thomas Maag
- No 5 Essays in Political Economy and Fiscal Policy**
Frank S. Somogyi
- No 6 Modern Statistical Methods Applied to Economic Time Series**
Simone G. Elmer
- No 7 Analyzing Capital Investment and Unemployment Using Modern Statistical Methods**
Rolf M. Schenker
- No 8 Capacity Utilisation and Macro Dynamics from a Micro Perspective**
Eva M. Köberl
- No 9 Essays on the Production and Measurement of Knowledge Capital**
Thomas Bolli
- No 10 Deviations from Rationality – Essays on Inflation Perceptions and Expectations**
Lena C. R. Dräger
- No 11 Determinants of Training and the Performance of New Firms**
Tobias Stucki
- No 12 Measuring Media Sentiment – Essays on Its Impact on the Economy and the Financial Markets**
Matthias W. Uhl
- No 13 Rigidities, Institutions and the Macroeconomy**
Heiner Mikosch
- No 14 Applying Modern Econometric Methods to Swiss Unemployment and Consumption**
Martin Straub
- No 15 Determinants of Cross-Border Activity – Essays on Trade and Entrepreneurship**
Andrea T. Lassmann
- No 16 Three Essays on International Trade Policy and Welfare**
Sergey Nigai
- No 17 The Generation and Diffusion of Energy Efficient Technologies in Switzerland: An Empirical Study**
Marius Ley
- No 18 Four Essays on the Economics of Crime**
Marie Poprawe
- No 19 Examining the Special Case – Special Trends and Events at the Swiss Labor Market**
Michael Siegenthaler
- No 20 Empirical Studies on Research Productivity, Labour Mobility, and the Swiss Hospitality Industry**
Jörg Schläpfer
- No 21 The Globalization of Goods and Capital Flows**
Benedikt Rydzek
- No 22 Taxation and Firms' Investment Decisions**
Simon Bösenberg
- No 23 Essays on Applied Nonparametric Econometrics**
Sabrina Maria Dorn
- No 24 Four Essays on the Estimation of Time-Invariant Variables in Panel Data Models**
Michaela Kesina
- No 25 Survey Based Research in Economics – Essays on Methodology, Economic Applications and Long Term Processing of Economic Survey Data**
Matthias Bannert

KOF

ETH Zurich
KOF Swiss Economic Institute
LEE G 116
Leonhardstrasse 21
8092 Zurich, Switzerland

Phone +41 44 632 42 39
Fax +41 44 632 12 18
kof@kof.ethz.ch
www.kof.ethz.ch
#KOFETH

