UNDERSTANDING SAVINGS BEHAVIOR:
A NEW SAVINGS TYPOLOGY BASED ON
INDIVIDUAL VOLUNTARY SAVINGS ACCOUNTS

A Thesis

Presented to the Faculty

of

The Fletcher School of Law and Diplomacy

by

ASHIRUL AMIN

In partial fulfillment of the requirements for the

Degree of Doctor of Philosophy

November 21, 2015

Dissertation Committee

JENNY AKER, Chair

CAROLYN GIDEON

KIM WILSON

DARYL COLLINS

This page intentionally left blank.

Ashirul Amin

Education

The Fletcher School of Law and Diplomacy, Tufts University Medford, MA
Doctor of Philosophy (PhD) Candidate 2010 – 2015

- Research Area of Interest: Forensic Analytics of Microsavings Portfolios in Microfinance Institutions
- Doctoral Fellow, The Hitachi Center for Technology and International Affairs (HCTIA)
- Provost Fellow

The Fletcher School of Law and Diplomacy, Tufts University Medford, MA
Master of Arts in Law and Diplomacy (MALD) 2005 – 2007

- Fields of Study: Development Economics, International Information & Communication
- Thesis – *From Village Phones to Telecentres: Developing an Explanatory Framework for ICT4D Tools in BOP Markets*

Princeton University Princeton, NJ
Bachelor of Science and Engineering (B.S.E.) in Computer Science 2000 – 2004

- Certificates in Near Eastern Studies, and Robotics and Intelligent Systems

Work Experience

Center for Emerging Markets Enterprises, The Fletcher School, Tufts Univ. Somerville, MA
Research Associate 2012 – present

Feinstein International Center, Tufts University Somerville, MA
Research Associate 2011

Bankable Frontier Associates (BFA) Medford, MA
Consultant 2011

Shakti Foundation for Disadvantaged Women (SFDW) Dhaka, Bangladesh
Consultant 2010, 2012

TriLinc Global, LLC. Manhattan Beach, LA
Chief Software Architect October 2009 – June 2010

MF Analytics, Ltd. Cambridge, MA
Associate 2007 –2008
Chief Software Architect, PAES 2008 –2009

Development Research Network (D.Net) Dhaka, Bangladesh
International Intern Summer 2006

You Can Hear Me Now Medford, MA
Marketing Assistant Fall 2006 – 2007

ALK Technologies Princeton, NJ
Developer/Analyst 2004-2005

Service Experience

Fletcher School Situation Room, Haiti Earthquake Medford, MA
Volunteer and Project Lead for Crisis Camp January 2010

Open Source Computer Education Program (OSCEP) Medford, MA and Dhaka, Bangladesh
Program Manager 2007-2011

Publications

Amin, Ashirul. "The Challenge of Combining Credit and Savings: (Innovations Case Discussion: Jipange KuSave)." Edited by Philip E. Auerswald and Iqbal Z. Quadir. *Innovations: Technology, Governance, Globalization* (MIT Press) 6, no. 4 (2011): 43-48.

Amin, Ashirul, Ahmed Dermish, Denise Dias, and Matt Herbert. *Is Grandma Ready for This? Mexico Kills Cash-Based Pensions and Welfare by 2012.* Case Study, Medford: Center for Emerging Markets Enterprises, The Fletcher School, Tufts University, 2011.

Amin, Ashirul, and Robert Mora. *Portfolio Analytics and Risk Management Practices of MFIs: A Global Survey*. Washington, DC: USAID, 2011.

Miscellaneous

- Editor: Center for Emerging Markets Enterprises Blog
- Computer Skills:
  - Programming languages: C, C++, Python, Perl, PHP, Java, Java Script
  - Database platforms: DB2, MSSQL, MySQL
  - Mathematical/Statistical software: Stata, R, SPSS

Abstract

Mass-market retail banks targeting low-income individuals offer no-frills voluntary accounts to make banking services accessible to the unbanked or underbanked. The granular transaction and balance information collected for each client is a treasure trove of behavioral data, yet a lack of analytical frameworks, computational power and motivation makes them difficult to analyze.

To enable us to delve into such large datasets better, this thesis develops a savings typology discovery methodology that relies on finding patterns in episodic periods of savings (or “segments”). It harmonizes each segment to make them comparable to each other, uses k-median cluster analysis to bucket similar patterns, and classifies recurring behavioral savings patterns, or “motifs,” into five groups: Accumulators, Sustained Balances, Fast Drawdown, Slow Drawdown, and Dump-and-Pull.

The Dump-and-Pull motif involves a short-term, pass-through mechanism that is not savings oriented. Accumulators are accretive savers who demonstrate that discipline-intensive behaviors are possible using bank accounts. The other three represent different levels of initial cashflow intermediation. Together, they show that the use of no-frills savings accounts do seem to further the cause of financial inclusion by accommodating a wide variety of behavior, savings and otherwise.

We evaluate the efficacy of motifs by exploring the additional explanatory capacity they provide in understanding banking agent usage, deployment of which occurred during the period covered in the dataset. We utilize Arellano-Bond GMM Estimators to explore the relationship between agents, motifs and outcome variables, with a special interest in the interaction between agent usage and motifs. We find that Accumulators and Sustained Balance motifs show increased interaction in smaller amounts, Fast and Slow Drawdowns have more lumpy interactions, and Dump-and-Pull show an overall decrease in involvement. The primary association with agents is through changes in the number of deposits, and then, by how much is withdrawn.

This page intentionally left blank.

Dedication

To my parents and family, for their endless love, support and encouragement!

This page intentionally left blank.

Acknowledgements

This dissertation would not have been possible without the help and support of many individuals. I am extremely grateful to them.

First and foremost, I would like to express my very great appreciation to the highly esteemed dissertation committee members, Prof. Jenny Aker (Chair), Prof. Carolyn Gideon, Prof. Kim Wilson, and Daryl Collins. Between Kim's invaluable mentorship for almost a decade now, Jenny and Carolyn's meticulous attention to detail and rigor, and Daryl's unprecedented experience in applied research in the field of financial inclusion, this thesis has benefited incredibly from all their expertise and wisdom over these years.

I would like to extend a special thanks to Elke Jahns, whose advice and inspirations at crucial moments of this process proved invaluable to keeping this project on track and in focus, and to Jenifer Burckett-Picker, for her encouragements to keep things moving when I needed it most.

I wish to acknowledge the assistance of Prof. Ehsan Hoque at the Dept. of Computer Science at the University of Rochester in ascertaining the veracity of the computer code used in parts of this thesis.

I would like to express my gratitude to the Kenya bank which kindly gave permission to use an anonymized sample of the bank's dataset, without which this research would not have been possible, and to Bankable Frontier Associates (BFA) for making that dataset available.

Last but not least, I would like to thank my family, friends and colleagues for continuing to support and encourage me through this long and arduous process.

This page intentionally left blank.

Table of Contents

Abstract v

Dedication vii

Acknowledgements ix

Executive Summary 1

Introduction 4

Chapter 1: Literature Review 9

The Poor and Their Savings 9

The Evolving Universe of Agent Banking 29

Reflections on the Literature 45

Chapter 2: Discovering Motifs 50

Overview 50

The Savings Account Dataset 51

Discovery of Account Usage Motifs 54

Chapter 3: The Agent Difference 93

Identifying the Intervention 93

Time Bound Comparisons of Means 94

The Need to Delve Further 113

Arellano Bond GMM Estimators 119

Conclusion 168

Key Findings 168

Implications of Findings 174

Further Research 183

Appendix A. Calinski–Harabasz pseudo-F Index Scores 185

Appendix B. Detailed Results of AB-GMM Runs 186

Bibliography 202

This page intentionally left blank.

Table of Figures

Figure 1. Mobile Money User Segmentation 41

Figure 2. Fungibility of the Concept of SAving Up/Down/Through 45

Figure 3. Overview of Motif Identification and Evaluation Process 50

Figure 4. Four sample accounts with withdrawals (red), deposits (green) and balance (blue) 57

Figure 5. Details of the PIP procedure. 62

Figure 6. Segments created by progressive cleaving at 10% threshold. 65

Figure 7. Splitting Balances into Equal sub-segments 71

Figure 8 progressive cleavage of balance segments to arrive at 5 balance values. 72

Figure 9. Progressive Merge of Balance Segments Down To Five Balance Values. 73

Figure 10. Harmonizing Balance Values to Percentages 73

Figure 11. Calinski–Harabasz pseudo-F Index Scores, by Cluster 76

Figure 12. Pattern Clusters and their Calinski–Harabasz pseudo-F Index Scores. 77

Figure 13. Details of Each of Six Clusters/Motifs. 79

Figure 14. Segments Matched to Motifs 82

Figure 15. Transaction and Balance Figures for Motifs 83

Figure 16. Avg Balance by Deposit and Withdrawal Frequency 85

Figure 17. Average Balance by Deposit and Withdrawal Frequency, by Motif 86

Figure 18. Motifs Assigned for Account (a). 87

Figure 19. Motifs Assigned for Account (b). 88

Figure 20. Motifs Assigned for Account (c). 88

Figure 21. Motifs Assigned for Account (d). 89

Figure 22. Balances Before (Left) and After (Right) Agent Usage 96

Figure 23. Natural Logarithm of Balances Before (Left) and After (Right) Agent Usage...... 96

Figure 24. Log Transformation of Number of Deposits and Withdrawals 101

Figure 25. Log Transformation of Deposit and Withdrawal Average Amounts 106

Figure 26. Coefficients with 95% CI for Average Balance 129

Figure 27. Error Plots, Average Balance 132

Figure 28. Residuals plotted against second lag of log(balance) 132

Figure 29. Residual Plots by Agent Usage 133

Figure 30. Residual Plots against second Lagged Motif 134

Figure 31. Coefficients with 95% CI for Deposit Frequency 138

Figure 32. Error Plots, Deposit Frequency 141

Figure 33. Residuals plotted against second lag of log(Number of Deposits) 141

Figure 34. Residual Plots by Agent Usage 142

Figure 35. Residual Plots against second Lagged Motif 143

Figure 36. Coefficients with 95% CI for Avg Deposit Amt 144

Figure 37. Error Plots, Average Deposit Amount 147

Figure 38. Residuals plotted against second lag of log(Avearge Deposit Amount) 148

Figure 39. Residual Plots by Agent Usage 148

Figure 40. Residual Plots against second Lagged Motif 149

Figure 41. Coefficients with 95% CI for Withdrawal Frequency 151

Figure 42. Error Plots, Withdrawal Frequency 154

Figure 43. Residuals plotted against second lag of log(Number of Withdrawals) 154

Figure 44. Residual Plots by Agent Usage 155

Figure 45. Residual Plots against second Lagged Motif 156

Figure 46. Coefficients with 95% CI for Avg Withrawal Amt 157

Figure 47. Error Plots, Withdrawal Amount 160

Figure 48. Residuals plotted against second lag of log(Average Withdrawal Amount) 160

Figure 49. Residual Plots by Agent Usage 161

Figure 50. Residual Plots against second Lagged Motif 162

Table of Tables

Table 1. GAFIS Savings Typology 21

Table 2. OSA Average Balance Distribution 54

Table 3. Number of Segments per Account 66

Table 4. Number of Transactions per Segment 67

Table 5. Length of Segments (Months) 68

Table 6. Distribution of Segment Balances, Maximum and Average 68

Table 7. Modified Assignment of Motifs 80

Table 8. Motif Distribution for Illsutrative Accounts 90

Table 9. Time-dummy regression results for average balance 97

Table 10. Time-dummy Regression Results for Average Balance and Dominant Motifs 98

Table 11. Time-dummy Regression Results for log(Number of Deposits) 102

Table 12. Time-dummy Regression Results for log(Number of Withdrawals) 102

Table 13. Time-dummy Regression Results for log(Deposit Counts), by Dominant Motifs 103

Table 14. Time-dummy Regression Results for log(Withdrawal Counts), by Dominant Motifs 104

Table 15. Time-dummy Regression Results for log(Average Deposit Amount) 107

Table 16. Time-dummy Regression Results for log(Average Withdrawal Amount) 107

Table 17. Average Deposit Amounts after Agent Usage, by Motif Dominance. 108

Table 18. Average Withdrawal Amounts after Agent Usage, by Motif Dominance. 109

Table 19. Change in Balance and Transaction Metrics after Agent Usage 110

Table 20. Hausman test for Fixed vs Random Effects 117

Table 21. Arellano-Bond test for Autocorrelation 119

Table 22. System GMM Estimators for Average Balance 128
Table 23. Pairwise Comparison of Interaction Coefficients, Avg. Balance 131
Table 24. Omitted Variable Bias Reference Table 135
Table 25. System GMM Estimators for Frequency and Amount of Deposits 137
Table 26. Pairwise Comparison of Interaction Coefficients, Deposit Frequency 140
Table 27. Pairwise Comparison of Interaction Coefficients, Avg Deposit Amt 146
Table 28. System GMM Estimators for Frequency and Amount of Withdrawals 151
Table 29. Pairwise Comparison of Interaction Coefficients, Withdrawal Frequency 153
Table 30. Pairwise Comparison of Interaction Coefficients, Avg Withdrawal Amt 159
Table 31. Summary Interpretation of AB-GMM Estimators (↑ indicates positive bias) 163
Table 32. Direction of Agent-Motif Interaction Effect, Stat. Sig. Coefficients Only 164
Table 33. Direction of Agent-Motif Interaction Effect, All Coefficients 166
Table 34. Three Key Elements of a Savings Account Business Case 179
Table 35. Summary balance and transaction frequency data by motif. 180
Table 36. Contribution Margin by Motif 181
Table 37. Raw and Normalized Contribution Margin, by Motif 182

Executive Summary

The poor can and do save, utilizing a range of savings devices. Individual voluntary savings accounts have been an increasingly important addition to their repertoire of available instruments. However, both existing literature and techniques therein are inadequate to satisfactorily understand or analyze saving behavior that manifests itself in these accounts which house savings of the poor, are expensive for financial institutions to support, and whose utility policy makers are unsure of. This dissertation will **develop a new savings typology** by undertaking portfolio-wide account-level examination of saver behavior over multiple years, and assess its potential utility by using it to **explore behavioral changes associated with banking agent usage**.

We seek to identify recurring behavioral patterns, defined as motifs, through rigorous quantitative treatment of transactional and balance data. The dataset consists of about seventy thousand no-frills savings accounts from a mass-market bank in Kenya and contains granular transaction and balance information for 30 months. These no-frills accounts have few impediments to individuals using them as they see fit, and are designed with low-income individuals and the unbanked in mind. Balance profiles are split into segments that capture a full savings cycle of balances rise as it from and going back to zero, harmonized to allow comparison across varying amounts and durations, and clustered into five motifs that we contend captures prototypical savings behavior. These five motifs are characterized as:

- Accumulators: Small amounts are saved over time before the entire amount is withdrawn – these are the quintessential accretive savers,
- Sustained Balances: A certain level of balance is maintained for a significant duration of the segment before withdrawal,

- Fast Drawdown: The majority of an initial deposit is withdrawn very soon after and a residual balance is maintained,
- Slow Drawdown: An initial deposit is drawn down through regular, fairly even withdrawals,
- Dump-and-Pull: A deposit quickly followed by withdrawal of the entire amount, suggesting an opportunistic use of the account as a pass-through in a distinctly non-saving behavior

We evaluate the efficacy of motifs by exploring the additional explanatory capacity this framework provides in understanding the impact of banking agent usage. Agency banking represents the use of retail outlets by conventional banks to offer banking and payment services, leveraging mobile technology and point-of-sales devices, and has high expectations but also some concerns regarding its impact on saving behavior of the unbanked or underbanked. Our dataset allows the study of the impact of agents on balances, deposits and withdrawals because agent deployment by the bank occurred during the period covered in the dataset.

We utilize Arellano-Bond GMM Estimators to explore the relationship between agents, motifs and outcome variables. The AB-GMM method accounts for outcome variables that rely on past realizations of itself (i.e. are "dynamic"), explanatory variables that are not strictly exogenous, fixed individual effects, autocorrelation within individual units, and works well for datasets which have a few time periods but many individual units. The regression specification includes lagged outcome variable, current motifs and lagged motifs as endogenous variables, agent usage as an exogenous variable, and an interaction term between agent usage and motifs that is also our primary element of interest.

When considered on its own, we find that agent usage is not associated with a change in balances, amount of deposits or number of withdrawals, but *is* associated with an *increased* number of deposits and *decreased* amount of withdrawals for Accumulator and Sustained Balance motifs, *decreased* number of deposits and *increased* amount of withdrawals for Fast and Slow Drawdown motifs, and *decreased* number of deposits and amount of withdrawals for Dump-and-Pull motifs. Accumulators and Sustained Balance motifs increase granularity of interaction with the savings accounts, Fast and Slow Drawdown motifs increase the amount of funds intermediated through these accounts, and Dump-and-Pull motifs reduce their overall engagement. Overall, the primary association with agents is therefore with changes in the number of deposits, and as a consequence, how much is withdrawn.

We conclude by noting that the use of no-frills savings accounts do seem to further the cause of financial inclusion by accommodating a wide variety of behavior within the same account, and allowing replication of preferred behavior such as those seen in Accumulator and Sustained Balance motifs. The associated impact of agents and the business case for Accumulator motifs is found to be the most compelling; corresponding bottom-line impacts of other motifs are more variable.

Introduction

The poor can and do save, and have varied liquidity and illiquidity preferences. Savings can be difficult because income streams are often low, irregular and uncertain. They utilize a range of savings options accordingly. Traditionally, such options have involved everything from short term arrangements such as saving under the mattress and reciprocal lending to friends and neighbors, to longer term ones, such as saving with a ROSCA or purchasing assets in the form cattle. Over the previous decade or so, financial institutions (FIs) have increasingly started offering savings services to low income clients, either directly through individual accounts or through accounts tied to an entire savings group. These accounts are becoming an ever larger part of the financial portfolios of the poor.

FIs serving more than a small number of clients tend to maintain a management information system (MIS) that has data on millions, if not billions, of transactions executed by anywhere between a few thousand to millions of savers. This is a treasure trove of information that captures detailed behavior of savers as they undertake financial intermediation either as a complement or a substitute to other instruments in their financial portfolio. A review of existing literature in section *Reflections on the Literature* will demonstrate that current methodologies used to study savings are insufficient to satisfactorily understand or fully analyze saver behavior as they manifest themselves in these accounts.

This dissertation will propose a new savings typology based on the study of individual voluntary savings accounts. Voluntary savings accounts are savings accounts where the customer is "not obliged to save as part of a contract for some other financial service" and can choose "whether or not to save, and the timing and amount of savings and withdrawals" (Rutherford 2005, 23). Such accounts may be available to groups and individuals; group-based voluntary accounts are not of interest to this thesis. Because these voluntary accounts are generally devoid of significant constraints by way of compulsory obligations or linked relationships with other (often credit) accounts, we can expect savers to express their preferred behavior in the absence of behavior-modifying incentives. The potential to provide unprecedented insight into how individuals, particularly those of low income, save is the primary motivator for this dissertation. Incidentally, this freeform nature also provides the most formidable challenge to classifying behavior within such accounts – there are no fundamental metrics such as interest rate or term to build off of, as is the case for microcredit. Identifying relevant metrics early on in this thesis will therefore be of great importance, as this will allow us to better classify behavioral patterns.

The general thrust of evaluating the implications of any behavioral patterns we discover and the impact associated with using agents is captured by the following questions:

- Do no-frills savings accounts further the cause of financial inclusion?
- Does agent usage encourage "desirable" behavior?
- Can these findings tell us anything about the feasibility of offering no-frills accounts?

This research comes at a time when there is a heightened sense of interest in savings for low-income clients. There's been a general recalibration of the various elements of microfinance, with microcredit no longer being seen as the silver bullet to eliminate poverty. Microsavings, along with microinsurance, conditional cash transfers (CCTs), mobile wallets etc., are increasingly being emphasized as anti-poverty mechanisms. Indeed, some contend that demand "among even the lowest level of the economically active poor for secure, convenient and appropriately designed financial savings services" is often more than for credit services (Robinson 2001, 19).

A more holistic approach to financial inclusion that sees microsavings as a critical component of the financial portfolio of the poor is largely a result of the new paradigm that has been embraced by the industry, when it came to be generally accepted in the 1980s that microfinance could "provide large-scale outreach profitably", as a result of which "microfinance began to develop as an industry" in the 1990s (Robinson 2001, 54). This new paradigm contends that not only do the poor not need to be taught to save because they already save in a variety of forms, but that there is, in fact, "massive demand for institutional voluntary savings among the poor". Indeed, FIs need to "learn in what forms and for what purposes the poor save, and to design instruments that meet the demand better than the savers can do by themselves". This has led to FIs offering products that attempt to provide "security, convenience, liquidity, confidentiality and returns", contributing to "financially sustainable institutions with high microfinance outreach" (Robinson 2001, 87). As we will see in Chapter 1, our understanding of what the poor do with these accounts is not nearly as developed as the intent and interest with which these products are distributed. Thus, we wonder, do no-frills savings accounts actually further the cause of financial inclusion?

Furthermore, the nature of the data available lends itself to a whole host of additional empirical analysis based on algorithmic techniques. The possibility of designing a toolkit that will allow financial institutions and other stakeholders in the inclusive finance industry to better understand, develop and monitor savings products is a strong secondary motivator. FIs do not undertake the level of analysis this dissertation will attempt because forensic portfolio analytics is a resource intensive endeavor, both in terms of computing resources and staff capabilities. For organizations that are often growing rapidly with stretched IT systems that are hard pressed to produce reliable accounting records regularly, let alone seek business intelligence on revenue generating credit-oriented products, taking deep analytical dives into savings portfolios are usually not too high on the priority list in terms of resource allocation.

Yet we know that at its most fundamental level, savings consists of deposits and withdrawals, the residual of which dictates the balance over a period of time. Its more structured sibling in the inclusive financial products arena, microcredit, is defined by a principal amounts, interest rates, and repayment terms, which in turn allow for derivative metrics such as duration, delinquency, prepayment rates etc. to be calculated. This allows credit-providing institutions to refine product offerings by tracking performance with arbitrary granularity. Is not a similar structured approach possible to analyzing savings? Identifying fundamental patterns of saving behavior by bringing a comparable level of analytical rigor to voluntary savings accounts feeds directly into the desire to develop a behavior-identifying methodology.

Agency banking represents the use of retail outlets by conventional banks to offer banking and payment services, often in the form of a corner store that people would already visit to purchase everyday essentials. These banking agents have high expectations to further financial inclusion but also some concerns regarding its impact on saving behavior of the unbanked or underbanked. We are fortunate because agent deployment by the bank that provided us with our dataset occurred during the period covered in the data. This not only allows us to study how balances and transaction patterns changed as agents were used, but also test whether motifs we discover show differentiated responses to it, signaling its utility in identifying divergent behavior from the portfolio average.

The uncertainty around the impact of agents comes from not knowing how it affects behaviors that are considered "desirable." For example, given that these are "savings" accounts, an increase in balances associated with agent usage would have been a desirable outcome. Policy makers and practitioners are worried that agents will have the opposite, detrimental impact to savings accumulation as they would make it easier for savers to withdraw their funds – concerns we look into in the section *Savings and Agent Banking*. All other things being equal, we would also like to see individuals having more granular control over their accounts as it allows greater flexibility in how individual funds are managed. There is a general consensus that agents should increase transactions, but it is not clear if this is conducive to savings accumulation. It is also quite possible that different motifs will show quite different responses to agents. These uncertainties on the impact of banking agents motivate our exploration of whether they promote "desirable" behaviors.

Once we have identified the motifs and change associated with agent usage, we ask what the business case implications are of the findings of this thesis. It is great if no-frills accounts promote financial inclusion, but since they are provided by for-profit financial institutions, this service must be justified with respect to its impact to the bottom line if we have any hope of seeing such solution scale. We use a costing framework and stylized figures from a project our bank was a part of to explore the business case implications. Specifically, we estimate the impact of different balance levels on the net interest income, and the deposit and withdrawal patterns on transaction activity contribution costs. This allows us to understand which behavioral motifs are friendlier to the bottom line than others.

To this end, this thesis is structured as follows. In *Chapter 1: Literature Review*, we catalog what we know about how low-income savers use savings instruments and what motivates them, and our current state of understanding of the impact of agent banking on savings behavior.

Chapter 2: Discovering Motifs details the motif discovery process that yields five distinct behavioral patterns. We assess the utility of these motifs in *Chapter 3: The Agent Difference* by detailing the additional explanatory prowess they provide to understand savings behavior. In the *Conclusion*, we summarize our findings, and answer the three questions laid out earlier about no-frills accounts and financial inclusion, agent usage and desirable behavior, and the feasibility of providing no-frills accounts.

Chapter 1: Literature Review

We review existing literature to understand possible drivers of savings behavior, anticipate behavioral patterns in the data, and demonstrate that current methodologies used to study savings are insufficient to satisfactorily understand or fully analyze saver behavior as they manifest themselves in individual voluntary savings accounts.

We then catalog the current agent banking landscape, explore our understanding of the motivations and dynamics of low-income users who use banking agents, and predict the impact of using agents on their savings behavior. We intend to demonstrate the utility of the typology developed in this thesis by offering insights above and beyond what we would expect from canvassing existing literature.

The Poor and Their Savings

One in four adults around the world had savings at a bank, credit union or microfinance institution in 2011. There is, however, much variation in the proportion of population banked – the figure ranges from "45% in high-income countries, to 24% in upper-middle-income countries, to 11% in lower-middle-income and low-income countries" (Demirguc-Kunt and Klapper 2013). Within these markets, a household's demand for a bank account depends on both "market and non-market factors" such as "price, illiteracy, ethno-religion, dependency ratio, employment, wealth status … and proximity to a bank" (Osei-Assibey 2009). What occurs when an individual ends up saving in such a savings account depends on a host of factors we shall explore in this section. We look at why the poor save, why they don't save, products that are designed address such savings needs, and the latest analytical tools used to understand savings behavior to inform our search for savings patterns. We present material relevant to Kenya to the greatest extent possible to provide the in-country context to our dataset.

A Note on Microsavings Reviews

Before we dive into the literature itself, it is worth nothing that micro-savings is studied by groups of individuals who come from rather different backgrounds and have different motivations behind their explorations of the field. As prominent development economist Jonathan Murdoch laments (Morduch 2008):

> The literature on household savings can be placed in three bins. First, writing by academic economists. Second, essays with practical advice by and for practitioners. And, third, perspectives from historians, anthropologists, sociologists, and unaffiliated development scholars. The literatures proceed independently. The academics rarely engage with the practitioners, and most have little to say about how to translate ideas into action. The practitioners, for their part, seldom step back to truly question key assumptions and amass rigorous evidence. Economists tend to be particularly insular within the academy, and the "non-economists" are insular in their own ways.

This will be reflected in the literature we review below, though more of the recent work seems o have come from the practitioners' camp than the others.

We will only review literature that specifically focuses on developing countries. The micro- and macro-environment that low-income households operate in are vastly different between developed and developing countries. As Prof. Angus Deaton at Princeton University notes, there are at least four reasons why we should study saving in developing countries as something distinct compared to saving behavior in developed countries (Deaton 1989):

- Demographic structure, household composition and income prospects are much different
- Lack of fiscal systems at the macroeconomic level that allow personal income manipulations

- General belief in postwar literature that savings is too low in these countries, and that development and growth are adversely affected by this
- Savings is difficult to measure – "data inadequacies are pervasive".

Why Do The Poor Save?

To state a truism, how the poor use their savings accounts will depend to a large extent on the motivations behind saving. Any typology developed must be able to incorporate the major motivations behind accumulating funds. We know that the poor often do not save by choice – they save because they must. The authors of *Portfolios of the Poor* note that a significant portion of the world's poor households live on less than $2 a day per head, but it is not the case that they actually have a steady income of $2 every day – far from it (Collins, et al. 2009, 2). Expenses such as food and school expenses are often accrued at predictable, if not regular, intervals. But the poor are often hit by a "triple whammy" when it comes to incomes – cash inflows are often low, irregular and unpredictable (Collins, et al. 2009, 35-40). The financial arrangements they have to undertake to match up misaligned incomes and expenses are considerable.

The poor also need to accumulate more than what each quanta of income brings in. As Start Rutherford puts it, "just because you are poor does not mean that all your expenditure will be in small sums"; indeed, "the poor need, surprisingly often, to spend large sums of money" (Rutherford, The Poor and Their Money 2000, 3). He breaks down the need for large sums of money into three categories (Rutherford, The Poor and Their Money 2000, 4):

- Life-cycle events, such as childbirth, marriage, building a home, widowhood, old-age and recurrent festivals. These events can be anticipated, even if their exact dates are unknown.

- Emergency needs, which can either be personal (e.g. sickness, injury or death) or impersonal (e.g. war, floods, fires etc.). These events require a “sudden and unanticipated need for a large sum of money” that is not available at home.
- Investment opportunities, which entail opportunities to spend large sums of cash by “investing in an existing or new business, or to buy land or other productive assets.”

Recent research corroborates this framework. An Accion study of five countries – Colombia, Dominican Repulic, Ecuador, Nicaragua and Peru – that looked into the financial behavior of rural residents found that there was an almost even split between those saving emergency funds and those saving with a goal in mind, with a small percentage waiting for a business opportunity. Half the savers saved in cash only, and a third of savers maintained a static savings amount, while saved funds grew over time for the other two-thirds. They also found that static savings are often “backup funds for emergencies, particularly health crises,” while the more goal-driven savings were incremental and family-oriented (Urquizo 2012).

In terms of sequencing savings with respect to expenses, the same report found that most savers (60%) save from what is left over after paying various dues, while fewer (12-15%) set aside savings before they start to spend, with the remainder saving when they have surplus income (Urquizo 2012). With respect to actual cash flows, a GSMA report on the Democratic Republic of Congo found that there is an almost even split between those who save every time they receive money (54%), and those who save on a set schedule (45%) (Gilman, Genova and Kaffenberger 2013).While the exact proportions may not be applicable in other markets, this does imply that savings for some will appear in their accounts when they are paid or receive a transfer, and before they have started spending, while for others, it will appear alongside payment of obligations when funds are deposited for that purpose, assuming the same account is used for such purposes.

Irrespective of the initial motivation to save, cash savings are the first mobilized when an emergency arises (Urquizo 2012). If the amount of cash savings is not adequate and a larger sum of money is required, the poor have three options – sell assets they own or expect to own, mortgage or pawn assets with the expectation to release those assets with future cash flows, or to find a way of "turning their many small savings into large lump sums" (Rutherford, The Poor and Their Money 2000, 5).

It is now generally accepted by academics and practitioners alike that "money management is, for the poor, a fundamental and well-understood part of everyday life" (Collins, et al. 2009, 3). Even the poorest of households will hold both debt and savings and use them in various combinations to meet financial needs. Collins et. al.'s Financial Diaries project found that of the 250 households it created detailed financial portfolios of, none used fewer than four types of financial instruments. Indeed, the average number of instruments used was about ten in Bangladesh, over eight in India, and ten in South Africa (Collins, et al. 2009, 15). Each instrument was also used multiple times. The "total cash turnover" compared to total monthly income was also quite high – 75%, 330% and 500% in households in Bangladesh, India and South Africa respectively (Collins, et al. 2009, 16).

A very recent financial diaries project from Kenya corroborates these findings, noting that the median household had ten sources of income reported in the survey year, and highlights significant levels of volatility such that "for the median household, income fluctuated ± 55% from month to month and consumption fluctuated ± 43%" (Zollman 2014). The poor thus seem to addressing the imperative to save by handling available instruments with considerable sophistication in the face of significant volatility of income and consumption.

This acceptance that the poor can and do save was brought on by the paradigmatic change that accompanied the commercialization of microfinance in the 80s and 90s we noted earlier. According to the old paradigm: a) the poor generally could not afford to save, did not trust banks and formal financial institutions and preferred to save in “nonfinancial forms” when they did save, and b) voluntary savings services could not possibly finance the FIs credit portfolio, and even when it did, defaults and losses were high. This resulted in a lack of savings mobilization, subsidized credit and a potential of high default rates risking the savings of poor clients (Robinson 2001, 86).

Having a multitude of options does not mean that the available combination of savings instruments is optimal. Many households keep more cash at home than they consider desirable because they don’t know what else to do with it (Robinson 2001, 235). The poor save through some combination of cash, grain and cash crops, animals, gold, silver, jewelry and other valuables, land, rotating savings and credit associations and regular savings and credit associations, raw materials and finished goods, construction materials, cash or grain lent out for profit, deposits with informal savings collectors, and labor obligations (Robinson 2001, 235). Different communities also display different degrees of affinity between saving and borrowing. In Kenya, low-income families display a greater affinity towards savings than borrowing, where the median household “held the equivalent of 129% of their monthly income in financial assets, versus the equivalent of about 53% of their monthly income in liabilities” (Zollman 2014).

Differences in livelihood within the same community also require the availability of different types of savings instruments. The five-country Accion study offers insights in this regard. Relatively few households combined income from both farms and microenterprises. Farmers find financial planning particularly difficult because they “cannot usually predict, let alone influence, the final sale income from crops,” leading a third of such households to “raise medium or large livestock for sale and small animals for personal consumption, but also as a marketable store of value.” Microenterprises were found in three flavors – commerce (65%), service (23%) and production (12%). Commerce microenterprises have “frequent small revenues and expenditures,” service enterprises have “less frequent large outflows and frequent small revenues,” and production microenterprises have “infrequent and irregular flows for both expenses and revenues”. These flows affect the "frequency of cash payments, deposits, withdrawals and need for loans," and by extension, savings (Urquizo 2012).

Having to juggle so many instruments, combined with irregular and unpredictable incomes, makes this an "intellectual and practical challenge" that the poor have to contend with every day – an insight that is often lost in the on-average "dollar-a-day perspective on global poverty" that otherwise does a good job on focuses attention on the fact that so much of the planet lives on so little" (Collins, et al. 2009, 17).

Gender is an important dimension when it comes to actualizing savings. Women have been reported to have "a more positive attitude towards saving" than men, and often save small amounts privately. Both men and women save, however, especially when it comes to household goals where larger amounts are involved. This is somewhat driven by the perception by women that "men can save larger amounts because they control the money and have higher incomes" (Urquizo 2012). For women participating in organized saving, such as through ASCAs, ROSCAs and saving groups, the support of their husbands is considered necessary for savings at the household level (Grameen Foundation 2013). Randomized Control Trials (RCTs) have found that women tend to benefit more than men because it "helped them gain greater control over the money, from their own temptation to spend it, and /or from requests for money from others". The study in Kenya found that savings accounts "appear to have increased women's ability to handle health emergencies," while the one in the Philippines study found that "savings increased women's economic empowerment" (Kendall 2010).

Thus, we find that the poor save for life-cycle events, emergences and in anticipation for investment opportunities. Such savings can happen with leftover funds or before spending starts, on a schedule or opportunistically, and cover a wide-ranging spectrum of amounts and frequencies.

Savings can be held in cash and in kind, are balanced with debt obligations and may be liquidated on short notice in an emergency.

Why Don't The Poor Save?

Not everyone is convinced that we can delve into the inner motivations of savers, however. Some caution that the poor may simply not save more because savings is too onerous a task, given subsistence conditions (V. Banerjee and Esther 2007, Rutherford 2005). Browning and Lusardi undertake an exhaustive review of major economic theories of savings and conclude that while current theories can accommodate various motivations to save, it is not clear that they can explain savings behavior (Browning and Lusardi 1996). Others pin the futility of this exercise on a belief that "the poor do not have a culture of saving and may prefer living one day at a time, with little planning for the future" (Bertrand, Mullainathan and Shafir 2006). Given the information presented in the previous section, we contend that it is highly unlikely the poor do not save not because they do not want to, but because they cannot. This section briefly explores some of those barriers that can get in the way of funds accumulation in savings accounts.

Not surprisingly, not having enough money is the most cited reason for not having a formal savings account, followed by that fact that many feel "banks or accounts are too expensive," and that someone else in the family already has an account that they have access to. Banks being too far away, not having the necessary documentation to open an account, "lack of trust in banks," and religious reservations are other significant reasons for not saving in bank accounts (Demirguc-Kunt and Klapper 2013).

Dupas and Robinson demonstrated that the barriers are not simply structural. In Kenya, they showed that despite waiving the account opening fee for a basic savings account, "63% of the people offered opened the account, but only 18% actively used it." Their survey evidence suggests that lack of trust in the bank, unreliability of service and prohibitively expensive withdrawal fees were the main reasons why the accounts were not used (Dupas, et al. 2012).

Behavioral biases can contribute to restrained saving behavior. Individuals can display "conflicted views about savings," where the benefits and advantages of savings are weighed within the context of a stated preference "for reinvesting all surpluses." Savings can be seen as "stalled funds, a luxury that they cannot afford," and as being easier than investing, which "requires work and involves risk" (Urquizo 2012). Priority is often given to repayment of debt when excess cash flow does occur. Households make trade-offs between short-term liquidity and long-term investments for the future. Even in cases where funds are earmarked as liquid savings, they are seldom idle funds. Rather, the funds are put to work to "provide immediate auxiliary benefits," perhaps by lending it to a neighbor through a ROSCA to help their business grow as our saver waits her turn (Zollman 2014). There is also a perception that if an emergency arises, the family will take care of things, rather than savings that may not even exist in sufficient quantities. If semi-liquid assets such as livestock are owned, it is assumed that those will be sold if a sudden need for cash arises (Urquizo 2012). Thus, savings can be seen as a sub-optimal choice in terms of utilization of funds, and as a bulwark against shocks.

Inadequate saving can also occur because individuals "tend to forget 'exceptional' (infrequent and relatively large) expenditure needs" (Karlan, Ratan and Zinman, Savings By and For the Poor: A Research Review and Agenda 2014). In so far as savings constitutes of foregoing current and certain consumption in favor of future and somewhat uncertain consumption, loss aversion may lead an individual to consume rather than save (Karlan and Morduch 2009). And finally, "information and knowledge gaps" can contribute to the ineffective adoption and use of savings products by the poor (Karlan, Ratan and Zinman, Savings By and For the Poor: A Research Review and Agenda 2014).

The poor, therefore, may not save in a savings account because it is expensive, inconvenient, unreliable, seen as suboptimal use of funds compared to investing, or inadequate compared to community support during an emergency.

Classifying Savings

Some savings devices have fairly well-defined structures in place, such as Accumulating Savings and Credit Associations (ASCAs) and Rotating Savings and Credit Associations (ROSCAs). In ROSCAs, members save the same amount in each meeting, and the "pot" is given out to one member every meeting, in turn. ASCAs differ in that they accumulate savings for a period of time before funds are distributed, often as a loan. Other savings devices, such as the piggy bank, are more freeform, where the quality and frequency of funds saved varies. Now that we have a sense of the many reasons behind why the poor save, or fail to do so, we turn to attempts to classify such behavior.

Existing classification systems for savings behavior can seem deceptively simple. One of the most widely referenced is SafeSave founder and veteran microfinance practitioner Stuart Rutherford's notion that savers either "save up", where they accumulate a lump-sum and then spend, "save down", where they essentially borrow the amount up front and then repay it back over time, or "save through", which is a combination of the two, as happens in savings clubs (Rutherford, The Poor and Their Money 2000). Irrespective of the "devices and services" used to save in each of the three options, the end goal for the poor saver is to accumulate a "usefully large lump sum" of money – a goal that also serves as a working definition of "savings".

Mark Schreiner, Director at Microfinance Risk Management, emphasizes the temporal nature of savings, defines it as "the movement of resources through time" and identifies three stages of saving: "putting in (depositing), keeping in (maintaining a balance), and taking out (withdrawing)" (Schreiner, Measuring Savings 2005). He contends that each of the phases has its own unique dynamics that must be studied in parts. Armendáriz and Morduch identify two types of savings behavior: "low frequency savings" associated with "steady, long-term accumulation" of assets, and "high-frequency savings", consisting of "short-term investments and … smoothing consumption". (Armendáriz de Aghion and Morduch 2005)

Some experts simply consider all of savings to be on other side of microfinance coin as microcredit, where the order cash inflow and the outflow are reversed. Collins et. al. call the loans "accelerators" and the savings, "accumulators", noting that both "help poor households … by exchanging usefully large sums for a series of small regular payments", making saving and borrowing "quite similar in practice" (Collins, et al. 2009, 110, 130).

The most elaborate typology to date has been suggested as a result of a recently concluded initiative called Gateway Financial Innovations for Savings (GAFIS), a project of Rockefeller Philanthropy Advisors, funded by the Bill & Melinda Gates Foundation (B&MGF), and managed by Bankable Frontier Associates (BFA). GAFIS worked with five banks in developing countries that offer savings accounts to low income clients: Standard Bank of South Africa, BANSEFI of Mexico, Bancolombia of Colombia, Equity Bank of Kenya, and ICICI Bank of India (BFA 2011).

One outcome of the project was a typology of savings behavior that looks at savings as interplay between liquidity, value and duration preferences, and classifies savings behavior into three parts (BFA 2011, 3):

- Type A: Low value, short term and completely liquid
- Type B: Some build-up of value, medium term, possible liquidity restrictions
- Type C: High value, long term, some liquidity restrictions

The study draws parallels to alternative savings possibilities by noting that Type A is akin to saving under a mattress, Type B is like saving in a savings club, and Type C is saving in a long-term asset, such as a cow (BFA 2011, 4-6). It also notes that there are patterns that can be discerned that involve financial intermediation but are not classified as savings. Two of these are the "dump and pull" behavior where a salary or government payment is deposited into an account and then completely (or almost completely) drawn down within a few days (called "Active but not Saving"), and where there are erratic deposits and withdrawals (BFA 2011, 6).

The residuals are classed as Balance Managers because their balance does not stay either above or below a threshold for multiple quarters at a time consistently, but bobs up and down across that level. They fall between those who show some consistent level of saving, and those that are simply using the account as a current account. We delve into the specifications of this classification system in more detail in the next section.

Measuring Savings

To compare and contrast the links between savings patterns and savings devices, we need to be able to measure savings behavior itself. The GAFIS project referenced earlier has the most elaborate quantitative enumeration mechanism publicly available. It formalizes the definitions based on two portfolio-based metrics – number of customer-initiated transactions, and average monthly balances (BFA 2012). In turn, it creates two indicators from these metric:

- Debit:credit ratio (i.e. number of withdrawals to deposits)
- Quarterly balance (derived from monthly balances) as a percentage of a minimum threshold

These indicators are then used to define the following categories (BFA 2012, 16-17):

Type	Number of transactions	Quarterly Balance	Note
Type A	Debit:credit at least 2:1	Balance every one of the quarters less than 50% of defined minimum threshold	Represents short term, low value, and fairly liquid savings, akin to saving under a mattress.
Type B	Credit:debit at least 6:1	Average annual balance is greater than first month balance	Some build-up of value, medium-term storage, akin to a savings club
Type C	None	Balance for three consecutive quarters is greater than 50% of the defined threshold minimum.	Higher value, longer term, akin to saving in physical assets.
Active but not Saving	Ratio of debits:credits less than 2 and credits:debits less than 6	Balance every one of quarter is less than 25% of the defined minimum threshold	No signs of savings accumulation, but active use of account.
Balance Managers	Residual	Residual	

Table 1. GAFIS Savings Typology

This represents a successful translation of a thematic framework into tangible parameters that allows any savings portfolio to be analyzed quantitatively.

A 2009 CGAP Technical Note by Joachim Bald offers a very detailed portfolio-centric analysis, looking at deposit portfolios in five institutions offering products for low-income clients. It analyzes long-term trends, core deposit trends, seasonal patterns, annualized daily volatility, and average life of demand deposits, as well as “peculiar patterns, trend breaks , and outlier values” (Bald

2008). The authors had access to monthly balance data over multiple years for all the banks, and monthly and daily balance data for three of the five (Abakaeva and Glisovic-Mezieres 2009).

Let us take a brief look at each of five quantifiable metrics, since it provides a look at the level of detail at which we have to conduct portfolio analytics:

- Long-term trend: Logarithmic regression is applied to monthly deposit balances, with dummy variables if necessary. The co-efficient provides an average long term growth factor (Bald 2008, 7).
- Core deposit trend: Local minima of balances are identified and an exponential curve is fitted to these minima that are considered to represent the "amount of long-term predictable funding generated by deposit-taking operations" (Bald 2008, 8)
- Seasonal patterns: A "relative index" is calculated that captures the "month-by-month" deviation from the long-term trend calculated above (Bald 2008, 8-9).
- Annualized daily volatility: The "standard deviation expressed in percent per annum calculated on frequent (ideally daily) logarithmic relative balance chance" is used as a measure of volatility, independent of the size of the portfolio (Bald 2008, 9).
- Average life of demand deposit: Understood as being the average time (calculated in days) that a unit of currency remains in an account, and calculated as follows (Bald 2008, 11):

Average Life = ∑ (End-of-month balance X No of days in month) / MAX(End-of-month balances)

Schreiner identifies seven measures of savings that can be derived from *monthly* deposits and withdrawals: gross deposits, gross withdrawals, participant accumulation, total accumulation, dollar months saved, dollar-months per month, and dollar-months saved ratio. (Schreiner, Measuring Savings 2005) He also defines "deposit frequency" and "deposit entropy" to measure deposit consistency. These measures are not reproduced in great detail because they are seen to be complementary to the definitions above.

One example of the use of savings portfolio data can be found in Schreiner's work, with Sherraden, that uses MIS data for Individual Development Accounts (IDA) in the American Dream Demonstration (ADD) program to show that the poor in the US can save by looking at crosstabs of savers with various characteristics. (Schreiner and Sherraden 2005) Note that this data "tracked monthly cash flows through IDAs owned by the poor in ADD".

Nevertheless, measuring savings behavior is difficult, both from a design and implementation perspective. Demand deposits require the "most sophisticated and costly management because [of] the volume and unpredictability of transactions" (CGAP 2003). Quantitative surveys involving financial information are tricky to work with because "people under or overstate financial transactions" and often confuse "stocks" of savings with "flows", thereby confusing the process with the result (Wright and Mutesasira 2001). In household surveys, savings "is not measured directly but is the residual between two large magnitudes, each measured with error," with one of the results being "household survey data often show an implausibly large fraction of households dissaving …". (Deaton 1989)

Part of the difficulty arises from the fact that money itself is highly fungible, it can be mediated through a large number of instruments, and there is an entire underlying socio-economic framework that dictates savings behavior. As Collins points out, "large, nationally representative economic surveys ... count the number of poor people worldwide and measure what they typically consume during a year .. but offer limited insight into how the poor actually live their lives week by week"; on the other hand, "anthropological studies and market surveys examine behavior more closely, but they seldom provide quantified evidence of tightly defined economic behavior over time." (Collins, et al. 2009, 3)

We end this section with an anecdote from Bald who conducted painstaking portfolio analytics that serves as a usefully cautious reminder of the limitations when engaging in an enterprise of comparable scale (Bald 2008, 2):

> The essential data requirement for this type of analysis is a long-run time series of tightly spaced aggregate deposit supply by product group (demand savings deposits, transaction accounts, and term deposits). Volatility of deposit supply and the resulting consequences for the liquidity of the institution manifest themselves on a daily basis, not on average monthly values. Therefore, daily (or at least weekly) data points are preferable for a meaningful analysis. Unfortunately, this type of data was surprisingly difficult for the institutions to reconstruct from their information technology systems. Although the study team would have preferred to conduct a much broader cross-sectional study on a statistically significant sample of deposit-taking MFIs and banks, the team had to be content with a handful of datasets from those institutions that invested the effort to generate the detailed data for the study.

Savings Products

We conclude our review of existing literature by looking at how our understanding of the poor and their savings needs has informed the design of savings products.

Voluntary savings products are categorized into some variation of the three buckets below, each satisfying a different kind of life cycle need (Rutherford 2005, 49, Hirschland 2005):

- Demand deposit accounts, also called passbook savings, for clients for whom liquidity is key
- Contractual savings accounts, for clients who seek to deposit small, often fixed, amounts of savings regularly to meet a “specific need at a specific point in time”

- Term deposits, for clients who forgo liquidity in favor of yield on a single large lump sum of funds by making a single deposit that cannot be withdrawn for a fixed amount of time.

The poor generally look for three things in an ideal savings product – frequency (including the ability to save daily), variability (to "cater to uneven, irregular and unreliable cash flows") and reliability (Hirschland 2005, 151). They most prefer demand deposit accounts because they do not require a regular income, and permit withdrawals at will – an option that is much appreciated even though it is not often exercised frequently (Hirschland 2005, 138). Contractual savings products have a market in small balance savers because they are easy to understand and "enforce the discipline needed to save for future needs", although the level of demand is not that high because "incomes may already be committed to loans or businesses that require regular payments" (Hirschland 2005, 139). Time deposits generally see almost no demand.

Savings accounts are often designed as commitment devices, where the commitment can be either "hard," or "soft." A hard commitment strategy involves tangible penalties, such as interest being forfeited for months when a deposit is not made, or "an agricultural savings account in which withdrawals before a pre-set target data corresponding with the sowing season incur a substantial penalty." A soft commitment approach is more psychological, where labelling an account as "school fees" makes the saver feel guilt or loss when funds are withdrawn for expenses unrelated to education (Karlan, Ratan and Zinman, Savings By and For the Poor: A Research Review and Agenda 2014).

Products are generally defined by a set of rules that determine what an account holder can and cannot do, with at least some of the rules being meant to operate as incentives to encourage a certain type of behavior. Wright notes the following list of product features (G. A. Wright 2005, 125):

- Opening balance requirement
- Minimum balance requirement
- Deposit minimums
- Withdrawal amount and frequency limits
- Requirement of notice for withdrawal
- Interest paid, including if it differs by balance amount
- Frequency at which interest is paid
- "Withdrawal, statement and ledger fees"

Various modifications to the rules are also offered based on customer and FI preferences (Hirschland 2005, 142):

- Time deposit or contractual accounts can provide a "stream of smaller payments" instead of one lump sum
- An event, such as a wedding, can trigger a payout
- Deposit and withdrawals can happen in kind, such as with grain
- Payments can go to a third party, such as "relatives, students, suppliers etc."
- Savings can be linked to a loan, such that "achieving a certain volume or term can trigger a loan that, when combined with saved funds, enables a purchase or other expenditure"

Note that while there are some completely freeform voluntary savings accounts, most will have a minimal level of rules. As long as they are not too onerous or enforce too much of a regimented behavior pattern, we can consider it to be a voluntary savings account. The specifics of the savings account we deal with are provided in the section *The Savings Account Dataset*.

Well-designed products "increase account balances by giving clients the option to withdraw their funds while motivating them not to [do so]" (Hirschland 2005, 152). Some incentives provided include tying interest rate and issuing lottery tickets to account size and automatically offering insurance above a certain balance. Some even use the minimum monthly balance as the metric to determine benefits instead of the average monthly balance, encouraging depositors to "not … withdraw funds without good reason" (Hirschland 2005, 152).

In addition to keeping funds safe, other important reasons to open savings accounts are "to establish relationships with the financial institution for subsequent access to credit," and "the need to facilitate payment transactions involving distance, including supplier payments and remittances" (Urquizo 2012). The most valued features of a savings program can depend on whether one already has access to such a product. A survey by Opportunity International Bank in Malawi found that the three most valued features for existing clients were "proximity to home/business," "security," and "low/no fee to open account," while for non-clients, it was "security," "high interest rate," and "good customer care" (Ferguson 2011).

Some believe that individuals primarily seek "services that allow for frequent small deposits and infrequent large withdrawals" (Ashraf, Karlan and Yin, Deposit Collectors 2006). Generally though, well-conceived products cater to the fact that most people want to save in a number of ways because their savings come from a variety of sources, such as "annual harvests, monthly remittances, … uneven profits of hawking, or withholding a handful of rice", which means that the poor want to "save single lump sum annually, smaller amounts weekly or monthly and very small amounts irregularly" (Hirschland 2005, 137). Despite this seeming range of products, FIs often fail to satisfactorily satisfy savings needs because their "frequency, amounts and terms of service are often too rigid" (Rutherford 2005, 24).

Inadequate product design can occur because microfinance in general is a "product-driven" business instead of a "market-driven" one, where the focus is on producing a good and then attempting to sell it, as opposed to identifying the needs of potential customers first and then designing products to fulfill that need (G. A. Wright 2005). Information opacity can also result in suboptimal products, as these markets are not the best understood and the costs of serving them are certainly higher than serving more affluent clients who maintain higher balances and who live in closer proximity in urban or semi-urban areas. As Hisrchland notes, "when serving small depositors, biggest challenge is usually not to design a new product that is unique to a specific market but to find an overlap with what people want and what the [FI] can manage cost effectively" (Hirschland 2005, 137).

The cost for inappropriate design is quite high. Ill-conceived products can be a result of the view that the poor will value whatever service the FI provides because they are the only available source of financial services for them – that the demand for such services is inelastic. This is almost never true, since the poor will have developed techniques to store money long before the arrival of any "formal financial service provider" (G. A. Wright 2005, 117). Ill-designed products can cause large drop-out rates amongst low-income clients – in East Africa, such rates are 25-60% (G. A. Wright 2005, 117-118).

From the service provider's point of view, it is sometimes suggested that the poor do not or will not save in banks, "aggregate value of savings is too small to be worth capturing in the formal sector" and "banks cannot collect small savings profitably" (Robinson 2001, 249). The cost of serving low-income clients is a particular concern because they "are thought to transact more frequently than account holders with larger balances" (Bald 2008, 1). Financial institutions are, understandably, driven by profitability concerns to a very large extent, and their bottom line is highly susceptible to transaction and balance behavior (BFA 2012). An inadequate understanding of the needs of the market, combined with an inadequate understanding of what clients are up to once they have been signed-up, can have severe ramifications on that bottom line.

The Evolving Universe of Agent Banking

To obtain the necessary context to assess the developed typology based on its ability to explain behavioral changes associated with agent usage better, we will now explore what we know to be the impact of agent banking on savings. In this section, we present the concept of agent banking, its footprint in Kenya and elsewhere, expected impacts on savings, and borrow some lessons from the closely related phenomenon of mobile wallets. The gist of our findings is that while account engagement is expected to increase with these accounts, it is unclear how savings levels are impacted.

Branchless Banking and Banking Agents

Branchless banking has been defined as "the delivery of financial services outside conventional bank branches using information and communications technologies and nonbank retail agents, for example, over card-based networks or with mobile phones" (Pickens, Porteus and Parker 2009). Commercial financial service providers "offer banking and payment services through postal and retail outlets, including grocery stores, pharmacies, seed and fertilizer retailers, and gas stations, among others" (Lyman, Ivatury and Staschen 2006). Such branchless banking services are provided through retail agents. Additional characteristics of branchless banking include reliance on technology to "identify customers and record transactions electronically," ability to "offer at least basic cash deposit and withdrawal, in addition to transactional or payment services," "backing of a government-recognized deposit-taking institution, such as a formally licensed bank," and availability during normal business hours (Ivatury and Mas 2008).

There are two models of branchless banking through retail agents – one led by banks, and the other, by nonbank commercial actors. In the bank-led model, "the bank develops financial products and services, but distributes them through retail agents who handle all or most customer interaction.

The bank is the ultimate provider of financial services and is the institution in which customers maintain accounts. … Retail agents have face-to-face interaction with customers and perform cash-in/cash-out functions, much as a branch-based teller would take deposits and process withdrawals. In some countries, retail agents also handle all account opening procedures and, in some cases, even identify and service loan customers" (Lyman, Ivatury and Staschen 2006). This bank-led model is applicable to our dataset.

The other model usually involves a mobile network operator offering deposit and withdrawal services to an e-wallet, which is not connected to a bank account, but is usually backed by an escrow account held in a banking institution on behalf of the telecom. The amount of money that can be stored in the e-wallet is fairly limited, as is the amount that can be transacted through it on any given day. Such limits are usually a result of prudential regulations that allow account openings with basic documentation, but are careful to prevent money laundering and other illicit activities. This model is not the generator of our dataset, but we will visit it in the section *Learning from Mobile Wallets* to borrow lessons from it as it has a lot of dynamics in common with agent banking.

Agent banking delivery channels come in three flavors (Ivatury and Mas 2008):

- POS-enabled bank agent: "Managed by a bank and uses a payment card to identify customers."
- Mobile phone-enabled agent: "Managed by a bank that uses a cell phone to identify customers."
- Bank-provided account linked to a mobile wallet: "Bank account that is linked to a mobile wallet. The bank does not manage the agent, and pays a fee to the telecom for deposits and withdrawals." Unlike the first two channels, it is not necessary to visit an agent for non-cash transactions such as transfers to other individuals or bill payments, but only when depositing or withdrawing funds.

Note that even though the third channel uses the mobile network backbone for transactions, all funds are linked to a regulated financial institution, which makes it part of the bank-led model. Our dataset was generated by costumers who had access to POS-enabled bank agents.

Banking agents have three main advantages compared to their brick-and-mortar counterparts that improve the customer experience while making business sense for the bank. The first is that they decongest branches. Agents can be seen as “human ATMs” who provide “greater customer convenience” by offering “more points, fewer queues, [and] more direct interaction with their money” (Mas and Siedek 2008). Complaints against branches include “inquisitive and overburdened staff, long queues, excessive documentation requirements, limited and unreliable ATM and other IT resources, insufficient support for illiterate and sub-literate customers, inadequate or confusing information regarding the various accounts and services available” (Tiwari, Dhawan, et al. 2011). Much of this can be ameliorated by moving the traffic away from congested branches, as agents have customer trust and are knowledgeable of customer usage and habits (Dolan 2009). Some retail agents are also available for extended hours compared a bank branch, allowing greater flexibility in terms of when individuals can choose to bank (Dean 2011).

The second advantage is that agents offer a cheap distribution strategy even if the focus is on existing markets (Infosys Finacle 2012). Global Savings Forum reports that setting up an agent costs two to four per cent of the cost of a branch cashier, such that even at maximum capacity, the fixed cost per transaction for a branch cashier is 78 cents, compared to 11 cents for a POS-enabled agent and 4 cents or less for a mobile-enabled agent or mobile wallet (Veniard 2010). The cost to install a banking agent in Pakistan is US$ 1,400, while that of a branch is thirty times as much, and the overhead on the agent is around $300, versus about $28,000 for a branch (Deloitte 2012). And in India, “it takes 2.5M rupees to open a branch in a rural/semi-urban location” while a banking agent incurs no initial investment (Ballem, et al. 2013). This lower cost structure allows banks to cater to accounts that were previously unprofitable, and potentially pass on savings to customers in the form of reduced fees.

The third advantage is that they allow targeting new customer segments, often by expanding on existing geographical coverage. “Piggybacking on existing retail infrastructure” allows a retail agent to operate in areas in which “transaction numbers and volume might be too low to support a full-fledged branch.” Lower operating expenses of agents enable to “cater to new customer segments that were previously not sufficiently economically active, for instance lower income customers in peri-urban areas” (Mas and Siedek 2008).

Agent banking can also be combined with other social welfare initiatives, such as channeling conditional cash transfers (CCTs) (Kenya, Colombia and Brazil), selling insurance products (India), and dispensing food vouchers (World Food Programme in southern Africa) (Oxford Policy Management 2011). Channeling CCTs through agents could have a positive impact on savings too, as the case that CCTs encourage savings behavior has been made often (Winkler 2014) (Zimmerman and Moury 2009).

The Financial Access Initiative notes that digital system needs five features for it to be of benefit to customers: network penetration, availability of cash in/outpoints, trust, acceptance as a store of value, and integration (FAI 2013). We see all except “trust” being addressed by the inherent setup of the agent banking model; we will explore how “trust” is reinforced by using community-based agents in the next section. Thus, agent banking seems to satisfy most of the requirements to be beneficial to customers, at least on principle.

Agent Banking in Kenya and Beyond

Increasing outreach of banking services has often necessitated the involvement of alternative delivery channels. Based on state-level analysis for the period 1961-2000, Burgess et. al. demonstrated that that rural branch expansion in India as part of its social banking program significantly lowered rural poverty (Burgess, Pande and Wong 2005). Given its explosive growth in multiple geographies and ability to thrive in underdeveloped financial markets, agent banking has the potential to bring about similar impact. As Microfinance Opportunities points out, "the promise of branchless banking for low income users/clients is real, if not realized" (M. Cohen 2013). We detail the current scale of agent banking in Kenya, our country of interest, and some other parts of the world to note its expansive reach and to provide context to the inquiry of this thesis. Millions of individuals are using banking agents; understanding how it impacts their savings behavior is probably of no marginal import.

The Finaccess National Survey of 2013 gives us the most recent picture of Kenya's financial inclusion landscape. A third of Kenyan adults seem to have bank accounts, with more male respondents (36%) having one, compared to females (23%). Involvement in informal savings groups is at a similar level, though higher for females (34%) than males (21%). Overall, the use of savings products has increased over the last decade. In 2006, 52% had used a savings product while 38% never had one, compared to 63.3% having a savings product, and 25.6% having never used one in 2013 (FSD Kenya 2013).

Branchless banking has greatly increased the reach of financial services in Kenya. For 76% of rural population, the nearest financial service provider is a mobile money agent – it takes less time to get to such an agent than a bank branch or to a bank agent. A mobile money agent is close enough to walk to for 57.4% of adults, while 21.8% have a bank agent within walking distance, and 10.7%, a bank branches. Two out of three adults are aware of banking agents, but only one in ten has actually ever used one (FSD Kenya 2013).

Kenya is most famous for Safaricom's M-Pesa mobile money initiative when it comes to furthering financial inclusion through non-traditional channels. The scale of banking agent deployments is in a similar order of magnitude as that of mobile money agents. In the time that Safaricom rolled out 40,000 mobile payment agents, 10 banks brought more than 10,600 bank agents online, most of which belonged to Equity Bank, and Kenya Commercial Bank (KCB) (Cracknell 2012). Compared to its international peers, the Kenyan banking agents are moderately active – a 2012 CGAP study revealed that there are 87 transactions per day at an agent in Kenya, compared to 157 in Brazil and 25 in India (Chen and Thoumoung 2012).

The Kenyan bank that is the source of our dataset, is considered to be "Kenya's most successful, and highly innovative mass retail financial institution, reaching over seven million customers across the group" (Cracknell 2012). They began agent rollout in 2010, are seeing 20% of all cash transactions happen through this channel already, and plan to set up 20,000 agents in a few years (Cracknell 2012). By the end of March 2013, over 2.3 million customers had registered for agency banking and around 80,000 savings related transactions (deposits and withdrawals) were being conducted each day at 6,892 outlets – all within 2.5 years of its inception (N. A. and Mishra 2013). Additional information is provided in the section *The Savings Account Dataset*.

Similar trends of expansive agent outreach are seen in other parts of the world. Latin America has been a pioneer of the agent banking model, with Brazil leading from the front. In 2010, before Kenya's agent banking scene had taken off, Brazil already had more than 150,000 agents at 10.45 per 10,000 adults. In 2011, agents represented the only source of financial services in a fifth of the municipalities in Mexico, a quarter of those in Colombia, and two-fifths of those in Brazil and Peru. The role of these agents in payments and transfers is unquestionable; the current challenge for regulators is to utilize agent banking as "an effective entry point into the formal financial sector" (Lee 2012).

In India, the Cashpor program, managed by the Grameen Foundation, enrolled 100,000 clients by June 2013, within 15 months of starting off. The average member has saved USD 7.50, with balances increasing at a rate of 15% per month (Shah, Ganesh and Agarwal 2013). India is also an example of additional financial inclusion synergies being possible by combining agent banking (or banking correspondents [BCs], as they are called in India) with existing savings practices. Self-help groups (SHGs) collect compulsory savings from twenty to thirty individuals, but do not always have an immediate use for it, or a safe place to store it. An agent can allow convenient deposits, which are longer term and in larger amounts compared to deposits the banking agent typically handles. Managing SHG deposits also helps Indian banks meet the central bank's priority sector engagement requirements, whereby every under-served rural area has a certain number of active savings accounts (Ballem, et al. 2013).

Nevertheless, despite the expanding footprint of banking agents, "few account holders currently report relying on bank agents (whether over the counter at a retail store or some other person associated with their bank) as their main mode of withdrawal or deposit." Countries that are farther along than others include Bangladesh, Laos, Nepal and the Philippines, where more than 10% of account holders report using bank agents (Demirguc-Kunt and Klapper 2013).

Savings and Agent Banking

There is not that much literature on the relationship between agent banking and saving, let alone robust studies, partly because it is a fairly new phenomenon in the world of financial inclusion and the focus has been more on payments and transfers. We explore how commentators anticipate changes to transaction and balance amounts, and impute savings changes from it.

The fee structure for using banking agents can influence client behavior, and such structures are usually based on a transaction-basis, where the client is charged for every transaction conducted, as opposed to a flat month rate or for free. Unlike bank accounts with ledger fees, clients only pay for transactions they conduct – a provision that individuals prefer compared to a flat fee. Agents are accommodating of more transactions as they receive a percentage of every fee as commission. By being closer to the client, agents "benefit from additional revenue associated with transactions acquired by the agent, such as person-to-person transactions and bill payments" which would at least substitute for a lot of the transactions that would have occurred at a branch, and "proximity may increase their willingness to pay for these services and increase the number of transactions conducted through the channel" (Veniard 2010). In some countries such as Peru, not only do banks not charge for transacting at agents at all, but also explicitly prohibit agents from charging customers either to ensure that agents remain the lowest cost channel (Reyes and Dias 2010).

While commissions earned are often the primary incentive, "increase in traffic and thus sales potential for agents who are retailers" works as a strong secondary motivator because it increases sales of whatever else they are selling in their store that is not agent banking related (Dolan 2009). In many communities, banking agents are also members of the community, who can help individuals who find handling money to be complicated, are illiterate and semi-literate, or are not comfortable dealing with staff at a bank branch. Clients note that "filling in withdrawal or deposit forms, opening a new account, or using new mobile technology would all be easier and faster with a familiar, patient" banking agent (Tiwari, Singh, et al. 2011). And when an emergency arises and money are required, a banking agent is close by to release funds and save one from borrowing from neighbors, relatives, or money lenders – arrangements that no one likes (Tiwari, Singh, et al. 2011). Thus a case could be made that agent banking encourages greater engagement than bank branches.

The case for higher balances is not so clear. On the one hand, clients are encouraged to "bring money into the system," since the bank can earn float income on the balance, and "by definition will bring in a second transaction fee associated with a withdrawal, transfer or a payment later on that can be split between the bank and the agent." In contrast, a cash-in intended for bill payments and loan repayments does not have a float income proposition for the bank, and the fee paid by the client is split three ways between the bank, the agent and the utility or loan provider (Mas and Siedek 2008).

From the clients' point of view, a banking agent offers the flexibility to save at the time and in the amounts, of their choosing, which could incentivize them "to make more frequent deposits, even if these are small in value." Reinforcing the notion of female savers as a distinct demographic introduced in the section *Why Do The Poor Save?*, female account holders note that they "like the possibility to a separate, and perhaps clandestine, account which they control, and to which their husbands do not have access" (Dhawan, Tiwari and Shukla 2011).

On the other hand, there are factors that could drive balances lower. Since the cash that clients bring in have to remain in the till of the agent, it increases the burden of handing more cash by the agent, since there are implications in terms of secure storage, limited use of extra liquidity and more frequent trips to the bank to deposit additional cash (Mas and Siedek 2008). In Brazil, for example, the incidence of paying bills using banking agents once a month is very high. These cash inflows far surpass any cash outflows happening during that period, causing even small retailers to accumulate up to US$ 20,000 in cash in one day. To mitigate risk, banks ask their agents to refuse larger payments or require them to deposit money in a bank branch when a certain limit is reached (Siedek 2008). Trips to the bank are fraught with danger, as agents may be robbed on their way to and from banks. Cash-in-transit insurance exists in some countries, but only partially covers such losses, with the banking agent having to "absorb the cost of the insurance and theft" (Tiwari, Singh, et al. 2011). In cases where agents are used for cash outflows, such as CCTs, there could the opposite effect, where the till runs dry as withdrawals are made, unless there is a specific infusion beforehand.

Coupled with the fact that transaction fees are often tiered, where smaller transaction amounts pay a higher percentage of the amount in fees, it is possible that a real disincentive exists to go beyond a certain cash ceiling at the tiller. As a result, a case could be made that given "lower transaction costs and a transaction driven revenue model rather than a float-driven one, agent banking systems are most cost effective for transactional accounts with low balances and frequent transactions" (Veniard 2010).

Even though agents have the potential to have a constructive role in saver-bank interactions, "basic business skills gaps, lack of customer trust, and limited ability to partner with large corporations" may act as impediments to letting them do so (Dolan 2009). Intermediating financial services is more time-consuming, and fraught with more risk than selling airtime for telecoms, which is often what these retain agents start out as. In India, MicroSave found that "persuading prospective customers, particularly poor ones, that recurring deposits, insurance premiums, or even mobile loan payments makes sense requires time, patience, sophisticated sales skills, and full bank support" – conditions that exist for few agents (Tiwari, Singh, et al. 2011). Not being able to provide the help when needed is a major contributor to poor customer service. Focus group participants in India, the Philippines and Zambia "expressed insecurity about both where to get help, and the cost of doing so when the office is far away," in cases where a visit to some office is required when the agent is unable to assist (M. Cohen 2013).

Between not having enough float and not being able to bridge client skill gaps, there is the potential to create a vicious cycle, where "unfulfilled expectations contribute to clients' mistrust." Market research reported by Microfinance Opportunities notes that "the branchless banking providers were not meeting some clients' expectations in part due to miscommunications or misunderstandings" (M. Cohen 2013). When expectations are not met, it is conceivable that engagement and funds trusted to agents fall.

Some savers might decide not to try out agents at all. Some demand free deposits *and* withdrawals, as is provided by a bank, and may prefer saving through post offices to receive the same benefits. For others, the trip to the town is not just about banking – they visit markets that are close to the bank branches for trade and groceries, enjoy the outing and do not want to give them up, and therefore don't see a reason to pay the agent for what is not seen as an inconvenience (Tiwari, Singh, et al. 2011). Despite how reliant many branchless banking initiatives are on technology, it is not clear that they all offer "good security guarantees to their beneficiaries." This is troubling "because of the massive cash flows these systems tend to generate, and at the same time, the limited educational background and negotiating power of the users they serve," making security issues of particular concern (Panjwani 2011).

All these risks need to be evaluated with context, specially comparing to existing alternatives. For example, "the majority of low-income Kenyans use informal methods to send money home," giving money to "friends and family members traveling back to the rural area." This is cheap but risky, as funds can get lost in transit. Funds are also transferred with considerable risk of loss "through bus and *matatu* (shared taxi) companies" that are not licensed to transfer money (Mas and Morawczynski, Designing Mobile Money Services: Lessons from M-PESA 2009). Compared to this, many prefer to pay the agent fees and have the funds be transferred electronically near the recipient's doorstep without any actual transit of physical cash.

Governments are lending their weight to agent banking too. Five years ago India's RBI introduced No Frills Accounts (NFAs). Of the approximately 50.6 million accounts that were opened, too many are used for withdrawing government benefits and wages only. The majority of accounts are inactive or dormant, and reportedly, 20% or fewer use their accounts for small savings. The RBI is now looking to banking agents (or BCs) to solve this inactivity and dormancy problem (Platt, et al. 2011).

Learning from Mobile Wallets

We end this chapter by taking a look at possible leads on client behavior from the world of mobile wallets, which we briefly introduced in the section *Branchless Banking and Banking Agents*. Mobile money and agent banking are similar in that they allow intermediation of funds electronically, without actually ever having to show up at a bank branch, and with cash-in and cash-outs taking place at designated agents. There are some differences between the two models. Mobile wallets offer the convenience of transferring to other users and pay bills directly from the mobile phone while an agent banking client needs to visit a physical agent for such transactions. Mobile wallets are limited in that they include a lower ceiling in terms of how can be stored in the wallet, no interest is earned on balances, and there is no access to other banking services that are available to an actual bank account. Despite these differences, we believe that there are enough similarities in benefits offered, client segment targeted and mode of operation that the lessons from mobile money could be relevant to agent banking to some extent – a case also made by Ivatury and Mas in their chronicles of the early days of branchless banking. (Ivatury and Mas 2008) .

The ubiquity of the mobile phone has catapulted mobile money transactions as a preferred means of handling funds for the unbanked in lower income countries around with world, to various degrees of success. In Kenya, for example, "43% of adults who report having used mobile money in the past 12 months do not have a formal account;" the figure is 92% in the Sudan (Demirguc-Kunt and Klapper 2013). As of 2013, there were 219 live deployments of mobile money services in 84 countries, up from 179 services in 75 counties in 2012 (Penicaud and Katakam 2014, 8). The Mobile Money for the Unbanked initiative at the GSMA offers the customer segmentation framework based on activity levels presented in Figure 1 (Levin and Camner 2013).

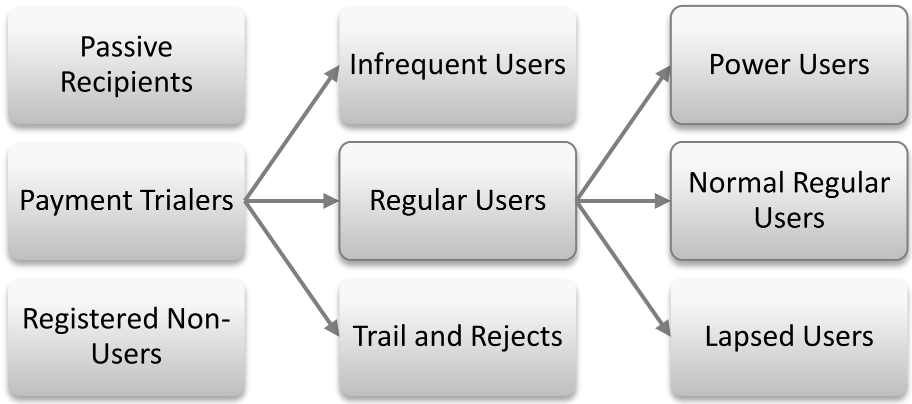

Figure 1. Mobile Money User Segmentation

According to this framework (Figure 1), registered users come in three varieties. “Registered non-users” register but never use the service. “Passive recipients” receive money from some source and withdraw the funds, but never initiate transactions themselves. “Payment trailers” initiate one or more transactions after registering. These “payment trailers” can become “regular users” or “infrequent users,” with the cutoff between the two being one payment a month. Some trailers can also become “trial and rejects” who never use the service after the first few tries. Finally, “Regular users” can be “power users,” who use the service more than five times per month and have more than five counterparties, “normal regular users,” who transact between once and five times a month and have at least two counterparties, or “lapsed users,” who used to be regular users but have not transacted in the previous three months. A similar framework is conceivable for banking agent users.

Investment payoff for the telecom usually occurs when users reach the “power users” and “normal regular users” stage only (boxes highlighted in Figure 1). In one particular study of a mobile money deployment, the same GSMA study found that only 18% of registered users became “regular users,” with a power users being only 4% of the user base, and normal regular users being another 9%. Thus, only 13% of clients made business sense for the mobile money service provider (Levin and

Camner 2013). While the proportions may vary, the gist of the findings is that there are few “power users” who transact at least once a week, and slightly more “normal regular users” who manage a few transactions a month.

Periodicity of mobile wallet transactions can be weekly, monthly or annual. A look into M-Pesa clients reveals that market days determine transaction volumes to a large extent, particularly in rural areas. Businesses are often closed on Sundays, leading to low transaction volumes. Rural towns with weekly markets have transactions concentrated on those days of the week. Monthly variations exist, with peak transactions happening “during the first week of the month, when salaries are typically paid.” Annual transaction patterns exist too, where transaction volumes decline in November and pick up again in December around Christmas, with a particular increase in withdrawals. Much of the savings in November is to allow individuals to save funds to bring to their families in person when they travel back home during Christmas. The study finds that the “variation from peak to trough can be as much as 40%, driving a wide variation in cash needs and store profits over the course of the month” (Eijkman, Kendall and Mas 2010).

There is some evidence that mobile wallets are used to save. Early reports seemed to indicate that saving is quite high with Kenya’s M-PESA users, with the percentage of users reporting “saving” up from 76% to 81% between 2008 and 2009, and those saving for emergencies from 12% to 22% in the same time period. This is mostly due to “early adopters saving more over time,” which suggests that “as users get familiar with the product, they are more likely to use it as a savings tool” (Radcliffe 2010). There is also some evidence that savings in mobile money can help cope with shocks. Jack and Suri finds that “while income shocks reduce per capital consumption by 7% for non-user households, the consumption of households with access is unaffected” (Jack and Suri 2014).

There is evidence to the contrary too, suggesting little or no savings takes place using mobile wallets. Analysis of data seems to contradict the propensity to save that respondents seem to recall.

The average stored value in an M-Pesa wallet is quite low (203 Ksh or US$ 3). The velocity of e-money, defined as the "frequency with which the average unit of money is used in transactions," is quite high at 11 – 14.6 transactions per month. The length of the e-money loop, defined as "the number of transfer transactions that the average unit of M-Pesa goes through between being transferred into a customer phone and being transferred back from a customer phone to the phone of an M-Pesa agent," is effectively one, suggesting that "the vast majority of transactions are a cash deposit, followed by a single person-to-person transfer, followed by a cash withdrawal" (Mbiti and Weil 2011). The same researchers subsequently found that "the average time a unit of M-Pesa remains on a user phone is about a week" (Mbiti and Weil 2013).

All this evidence in the Kenyan context points to the fact that mobile money is more transactional in nature, than savings oriented. A possible reason forwarded for low savings using mobile wallets is the lack of interest on balances. There is also evidence that "increased use of M-PESA lowers the propensity of people to use informal savings mechanisms such as ROSCAs but raises the probability of their being banked" but "little evidence that people use their M-Pesa accounts as a place to store wealth" (Mbiti and Weil 2011). Similar finds were reported in a study conducted on Orange Money customers in Madagascar in 2012 showed that "m-banking services increases the number of national remittances sent and received, but has no significant impact on the sums saved by users or the sums of remittances sent or received" (Arestoff and Venet 2013).

The recognition that it is difficult to save with mobile money is there in the industry, and some are rolling out "savings accounts that are distinct from mobile money accounts and which offer additional functionalities that are relevant for savings." Nine such services were launched around the world in 2013. Perhaps the most successful of these ventures is M-Shwari, a "credit and savings product" for M-PESA clients where Safaricom partnered with the Commercial Bank of Africa (CBA) to be able to offer interest on account balances. This allowed CBA to acquire 5 million additional deposit accounts, giving it the most customer accounts in Kenya after Equity Bank (Penicaud and Katakam 2014).

There is also a concern that mobile money might be loosening the mechanisms that traditionally encouraged savings. While the efficiency and low costs brought by mobile money may have enhanced financial freedoms of the hitherto unbanked, it does so by “challenging the social capital that many argue is key to development,” and involving strains such as “increased remittance solicitation and decreased rural visits, as well as to reductions in using friends and families as cash remittance intermediaries” (Donovan 2012).

Reflections on the Literature

On Savings Patterns

We find that while the existing literature offers some attempts at savings typologies, such as Rutherford's saving up/down/through, or Schreiner's putting in/keeping in/taking out triptychs, they are generally quite broad in their definitions and cannot be expected to adequately accommodate all the variations in behavior seen in these savings accounts. Consider a situation where a saver has been saving regularly for a long period of time, depositing \$10 every week and taking out \$50 every 6th week. Given that the cash flows are part of a continuum, whether the individual can be considered to be saving up, down or through really depends on the time envelope that we look at the cashflow patterns through (Figure 2):

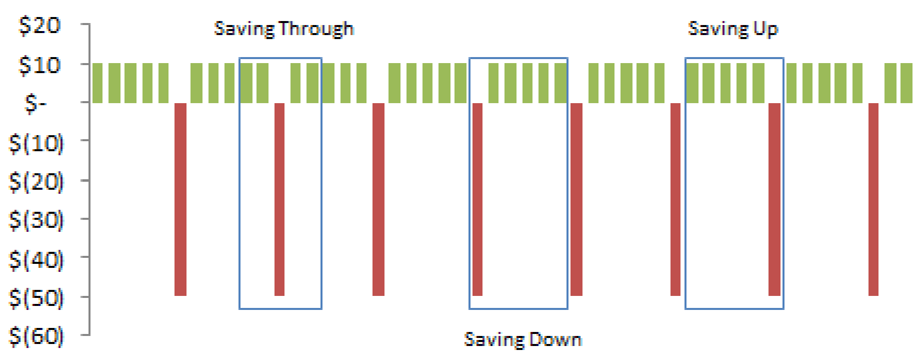

Figure 2. Fungibility of the Concept of SAving Up/Down/Through

Accounts where deposits and withdrawals can happen many, many times in the course of a year cannot be understood with a framework that inherently possess this level of ambiguity. Schreiner's triptych does not have this temporal uncertainty associated with it since it decomposes the process of savings into its three elements – deposit, withdrawal and maintenance of balance – but has nothing more to say in terms of the various combinations that these elements can exist in, which is key to understanding the enormous variety we see in savings behaviors.

BFA's Types A/B/C classification is a generation ahead in terms of sensitivity to different classes of behavior and explanatory ability compared to the more thematic frameworks, and does an admirable job in identifying accounts that mimic one of three dominant savings behaviors. Nevertheless, there are some potential shortcomings of this technique:

- The frequency of transaction types is not considered; instead a ratio is used. Individuals undertaking 2 withdrawals and 1 debit over a one year period will be lumped together with those who withdraw once a week while depositing once a month (at a ratio of 52:12) as Type A's, as long as their balances are below a certain threshold. It is reasonable to expect these individuals to be the same types of savers, or is frequency an important delineator of behavior?
- The amount of transaction amounts is not part of the equation. Yet an individual saving $1 at a time probably has different prerogatives than one saving $100 at a time. Do savers truly display the same patterns irrespective of how much is saved?
- Average balances over entire quarters are considered for balance benchmarks, which has a great smoothing effect. Yet many of these accounts display a wide range of volatility within a given year. Should such volatility (or lack thereof) not have be accounted for to completely capture savings behavior?
- The time period over which classification is done is fixed at a year. While this is reasonable to track seasonality variations that often have an annual cycle and provides some degree of protection from "noise", does this not loses out on capturing recurring patterns that manifest themselves monthly or even weekly?

The typology we develop will attempt to address as many of these issues as possible by incorporating adjustments directly into the pattern recognition process.

Incidentally, Bald does better in proposing sophisticated metrics, he does not attempt to combine them in various configurations to construct classifiers. In terms of setting standards for this thesis, Bald's work represents the baseline rigor in terms of analyzing relevant parameters, and A/B/C system sets the benchmark for effective juxtaposition of such parameters to construct a classification system that tries to organize entire savings portfolios.

In terms of what we can expect to see in the patterns, we should be prepared and attempt to account for any combination of balances, deposits and withdrawals. Balances could be accretive, when savings is planned, or static, in case it is funds earmarked for an emergency. Additions to the account could be regular, where money is put aside when wages or remittances come in, or irregular, where savings represent the occasional surplus after debt has been services and expenses paid. Similarly, savings draw-downs could be regular, when planned for periodic events such as school fees, or irregular, when an emergency arises or a plug for a business is needed. Deposit amounts could be large, perhaps from a crop sale, or small, as with microenterprises. So could withdrawals – we saw how commerce-oriented microenterprises have frequent, small expenditures while those for service-oriented ones are large but infrequent.

An underlying assumption behind this quest for a typology is that some of these variations around deposits, withdrawals and balance levels occur together more often than others. This notion makes intuitive sense and is reinforced by the classification frameworks we have visited earlier.

On the Impact of Agents

The second part of this thesis asks, once these savings accounts start using agents, do their behavior change, and does the typology developed help us understand that change better than what we would expect using existing literature? The gist of our expectations from literature is that the number of transactions will increase as savings account holders interact at agents, but we cannot predict whether savings balances will increase or decrease as a result.

First, the overwhelming consensus from the literature seems to be that agents will increase engagements of clients with their accounts. Shorter queues, closer proximity, friendlier staff and extended hours are all reasons to transact more, not less. Because of the lower cost structure of banking agents, banks can support accounts that were not sustainable before. And agents do not mind more transactions because they are paid on a per-transaction basis, unlike bank tellers. Thus, we can expect the number of transactions conducted to go up upon interaction with agents.

Transactions are essentially composed of inflows in the form of deposits, and outflows in the form of withdrawals and transfers. In so far as deposits are free at both agents and bank branches, there is no reason for these to go down in frequency. Withdrawals and transfers at agents have a fee, though usually at a lower rate than branches. Thus, there is also no reason to expect the number of withdrawals to go down as long as the fees are not higher at agents.

Second, the impact on actual savings levels is not clear. Savers can now deposit in small amounts with greater frequency in ways that would not have been feasible with a branch. The option to withdraw as needed and with haste is appreciated, even if most would not exercise that option except in emergencies. Agents can also be expected to encourage deposits to maintain cash in the retail till. However, if agents make it easier to deposit, they also make it easier to withdraw. The commitment necessary to build up savings, especially in the face of trying conditions, could crumble that much faster in the face of convenient withdrawal options. Agents may be reluctant to accept large deposits due to security concerns. And if expectations are not met or there is a breach of trust, savers may stop using agents altogether. We are thus unable to predict whether balances will go up, or down, after interaction with agents.

If savings accounts that operate through banking agents are anything like mobile wallets, we could expect to see most accounts exhibit low level activity. There would be low levels of savings but quite high velocity. We would also need to keep an eye out for weekly, monthly and annual periodicity.

Chapter 2: Discovering Motifs

Overview

The discovery and evaluation of recurring patterns in account usage behavior, or motifs, will follow the steps outlined in Figure 3 below. We focus on the three discovery-related steps in this chapter, and on the evaluation-oriented one in the next.

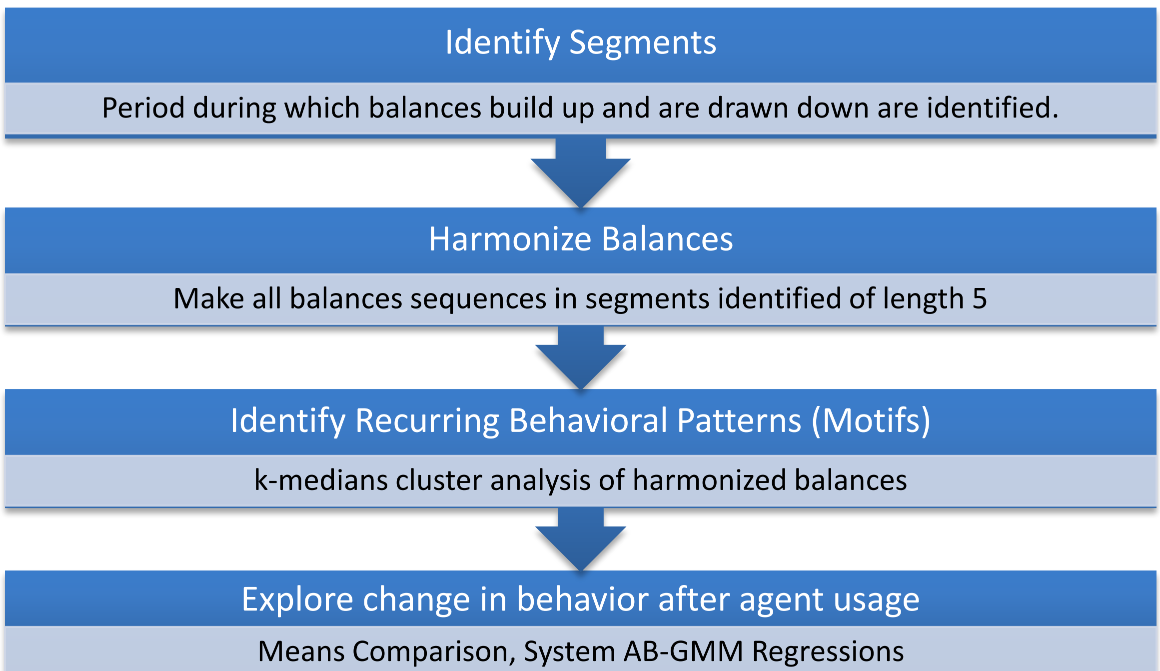

Figure 3. Overview of Motif Identification and Evaluation Process

The first step in identifying motifs involves demarking segments within which the balance of an account rises from and then falls back to zero. This is considered to be a complete savings cycle, and the behaviors expressed within these segments are our primary focus.

Segments can be quite different from each other, differing in balance amounts, days spanned, number of deposits and withdrawals, and amounts of those deposits and withdrawals. In the second step, we harmonize the amplitude of all segments by considering them as a percentage of the maximum value in the segment, and the length of all segments by transforming all of them to constitute of six sub-segments. If a segment has more than six balance values to begin with, we condense some; if they have fewer, the longest stretches of balances are split.

Finally, we apply cluster analysis to bag similar balance profiles, and arrive at five distinct motifs and one residual bucket. We apply a post clustering filter that weeds out segments that are weak matches, and assign them intuitive names of Accumulators, Fast Drawdowns, Slow Drawdowns, Sustained Balances and Dump-and-Pulls.

The Savings Account Dataset

Our dataset consists of transaction and balance data for 70,994 Ordinary Savings Accounts (OSA) from a leading retail bank in Kenya. They will be referenced henceforth as **Bank A**, which was a precondition for receiving this data. The data spans 30 months, from January 1, 2011 to June 30, 2013. It represents approximately 1% of all OSA accounts that were part of Bank A's portfolio as of June 30, 2013. This data was obtained courtesy of Bankable Frontier Associates (BFA), which had access to the entire dataset. BFA selected the random sample of one out of 10 accounts, and obfuscated the account IDs to make the dataset completely anonymous.

The transaction data is granular, in that every single transaction that occurred within those 30 months is available, along with a timestamp and a transaction type. The dataset consists of 5,077,207 such granular transaction records. The balance data consists of end-of-month balances for every account, for every month since inception.

Bank A describes itself as “the leading inclusive bank in Africa,” serving 9.2 million bank accounts which comprises of 50% of all bank accounts in Kenya (Source: Bank A material). They have a particular focus on affordable inclusive finance, a business model that “targets the low-income market to achieve scale through a high volume of relatively small, low-margin transactions,” and a “large distribution network of agents and a robust information technology platform further enable the bank to access previously untapped markets” (IGD 2013).

The Ordinary Savings Account (OSA) is Bank A’s flagship no-frills savings product, with “no ledger fees, maintenance fees, monthly charges or minimum operating balance,” or “cash deposit or cheque handling charges”. There are some charges, such as a KES 400 (~US$ 4.40) account opening fee, a KES 30 (~US$ 0.33) ATM fee for its own ATMs, and a counter cash withdrawal fee of KES 50 (US$ 0.56). These charges are significantly lower compared to its peers, and have fueled Bank A’s rapid growth in clients from 1.8 million in 2007 to 8.4 million in 2013. (Source: Bank A material.)

Bank A’s agent network has experienced incredible growth during the period we explore, with the number of agents increasing from 875 from the beginning of 2011 to 10,260 agents by December 2013. Agents accounted for a third of all cash transactions by the end of 2013, while the number of transactions at branches and agents remained fairly unchanged. Bank A credits banking agents for reduced cash handling costs and overcrowding in branches, and improved access to the unbanked (Source: Bank A material).

We consider Bank A’s OSA account to be an appropriate voluntary savings account to explore fairly unconstrained savings behavior as there are no restrictions imposed by the account in terms of how often or in what quantities transactions can take place or balances must be maintained. The handful of fees that do exist can be considered to be the cost of business with the bank by the clients. It is also fortuitous for us that we have data exactly from the period when agent-rollout began, allowing us a window into the account behavior before and after their existence.

Our interest in low-income savers is also well-served at a theoretical level by these accounts, as they are specifically targeted towards that client segment. Practically, however, since we do not have demographic data on clients, we cannot say with certainly what percentage of the portfolio are low-income. We offer circumstantial evidence that a significant proportion of Bank A's clients are low-income by comparing the account balances to income levels of our interested demographic segment.

For reference, a 2012 survey found that almost a third (29%) of Bank A's clients were poor when the $2/day income measure was used (BFA 2012). Amongst the low-income participants of the Kenyan Financial Diaries project that ended in 2014, we see that the average monthly rural and urban incomes levels are KES 1,706 and KES 4,800 respectively, while the average monthly rural and urban consumption levels are KES 1,824 and KES 3,651 respectively, compared to a national median monthly income of KES 2,167 (Zollman 2014).

We compare this with the account balance levels for the seventy thousand accounts in our sample. Table 2 gives us average balances for various percentiles for all accounts, and accounts that have greater than KES 0.00 in their account. Note that US$ 1 ~ KES 90, implying that 95% of the accounts have less than US$ 100 in their bank accounts. Table 2 also tells us that about 90% of accounts have less in balances than the average urban monthly income.

		5th	10th	25th	Median	75th	90th	95th	Mean
All Accounts (N = 70,994)	KES	0.00	0.00	0.00	32.40	436.73	2,468.72	5,836.43	1,603.58
	USD	0.00	0.00	0.00	0.36	4.85	27.43	64.85	17.82
Accounts w/ non-	KES	1.49	3.05	16.30	124.63	852.68	3,642.00	7,910.60	2,152.54

zero balance (N = 52,991)	USD	0.02	0.03	0.18	1.38	9.47	40.47	87.90	23.92

Table 2. OSA Average Balance Distribution

We emphasize that we cannot simply assume that all low balance accounts belong to low-income individuals, as middle- and high-income individuals could maintain low levels of funds in OSA accounts. However, given that the counterfactual to this situation is that seven million middle and high income savers opened OSA accounts yet decided to transact in relatively small amounts and leave small residual balances, we feel comfortable assuming that a significant proportion of accounts belong to low-income savers, especially when we combine it to deliberate targeting by the bank and recognition thereof through third-party publications.

Discovery of Account Usage Motifs

What Constitutes A Pattern?

In so far as motifs are essentially recurring patterns, it behooves us to define what constitutes a pattern first. For every account, we have the following:

- A record of transaction amounts over time. The amount can be either positive (credit) or negative (debit). Records are granular – one transaction constitutes one record, with no aggregation whatsoever by transaction types, or time.

- A record of balance amounts over time. Balances are reported at the end of day, for every day where there is a change in balance (i.e. one or more transactions have occurred).

The balance at the end of any given time window is equal to the balance at the beginning of the given time window, net of all debit and credit transaction amounts. In so far as the balance at the end of a period is the same as the balance at the beginning of the next period, the initial balance for the next period can be calculated by adding the sum of all credit amounts and subtracting all debit amounts of transactions occurring during a period to the starting balance for that period:

$$(\mathit{Initial\ Balance})_{t+1} = (\mathit{Initial\ Balance})_t + \sum (\mathit{Credit\ Transaction\ Amounts})_t - \sum (\mathit{Debit\ Transaction\ Amounts})_t \qquad (i)$$

This seemingly simple specification will be useful to draw inferences about matching behavioral patterns once motifs are identified. Note that for balances to match, we must include not only customer initiated (CI) transactions, but those generated by the business as customers interact with the bank. Such business initiated (BI) transactions include credits such as interest payments on account balances, and debits such as fees.

Balances capture the amount of funds accumulated in the account at any given time, and are intuitively what we relate to when we refer to "savings." This thesis therefore primarily focuses on balances. Theoretically, this thesis could also have focused on deposits and withdrawals, as how often the occur and in what amounts directly shape how balance patterns look over time, but as described in section The Savings Account Data, transaction timestamps contain the date but not the exact time. We can therefore calculate the end-of-day balance precisely, but there is no similarly meaningful end-of-day deposit-and-withdrawal construct. For example, three consecutive deposits of $1, followed by three consecutive withdrawals of $1 make for a very different pattern than alternative $1 deposits and withdrawals. However, the net balance at the end of the day for these six transactions is always $0.

The balance records consist of time series data for many accounts – panel data, by definition. Time series data typically have: a) a trend component that represents a long term trajectory that does not repeat over time (at least within our time window of interest), and b) a seasonality component, where some pattern does repeat itself over time. What is left over is the irregular component, or "noise".

Before we explore rigorous techniques to identify patterns, let us start with an intuitive sense of what we are attempting to accomplish. Figure 4 illustrates four accounts with fairly different account usage patterns, out of the 71,023 accounts that represent our sample. They are not necessarily representative of accounts in general; rather, they were cherry-picked because they allow us to "see" what is going on with them at a cursory glance. Accounts can demonstrate significantly different levels of transactional irregularity and balance level volatility.

Account (a)

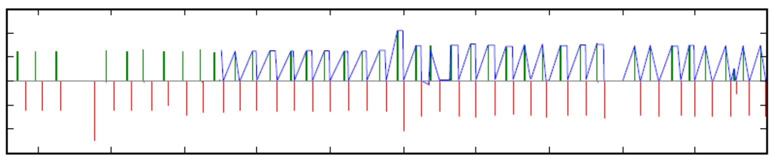

Highly regular behavior – deposits funds only to withdraw it all soon after. Seems to happen repeatedly, and it relatively fixed intervals. Often seen with salary accounts.

Account (b)

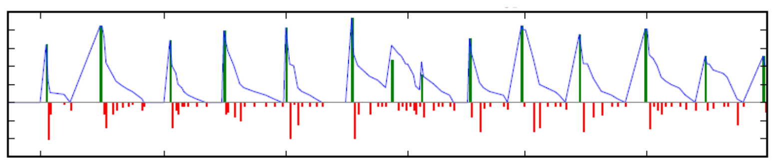

Fairly regular behavior – deposits funds and slowly depletes it over time. Multiple withdrawals represent the client's preference to retain at least some funds for a period of time. Perhaps monthly wages used as needed?

Account (c)

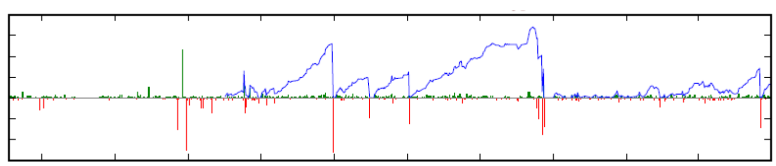

Some repeating behavior – saves small amounts over time, and withdraws everything as a lump-sum periodically. This requires diligence to build up savings, and is often associated with branch usage.

Account (d)

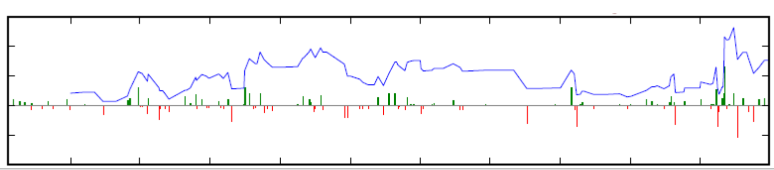

Balance manager – there is no clear pattern of what is going on, but the client is regularly using the account to intermediate funds. They deposit and withdraw somewhat irregularly, while maintaining a balance level. Account could belong to a microenterprise with much activity.

Figure 4. Four sample accounts with withdrawals (red), deposits (green) and balance (blue)

Accounts (a) and (b) show fairly regular behavioral patterns. While there is no long-term trend in the balances, (a) displays fixed periodicity, while the seasonality is less regular for (b). There is no long-term trend across all of (c), though arguably the gradient of the balance segments between drawdowns is similar. The periodicity for (c) would be difficult to pin down because the lag periods are different, the amplitude of the peaks are different, and the client seems to transition to a different usage pattern around the latter third of the time window before making a muffled attempt at accretive saving right near the end. (d) is devoid of any trend or seasonality that the naked eye can make out.

Before we can dive into classifying accounts, we need to decide what the unit of analysis is. This represents another intuitive leap of faith, but in the dance of deposits and withdrawals, one phenomenon we almost always see is that the balance will reach zero, or close to zero (especially when compared to the amounts transacted). We propose that **the balance profile contained within the time window between subsequent dips to zero (or near-zero) be treated as a "pattern"**. We will call each time window a "segment," where a segment is that part of the time series over which a pattern manifests itself. Thus, (c) above has four patterns in the middle of the time series, and a few more near the right.

Discerning behavioral patterns there involves identifying these structural breaks that define where segments begin and end, and then grouping similar patterns together in what is essentially a clustering exercise. Each cluster will contain the kernel of the typologies we're seeking. It is possible for an account to have more than one pattern, more than one type of pattern, or no identifiable pattern at all. We will formalize this clustering exercise in the next section.

Clustering Patterns

Cluster analysis is “the art of finding groups in data” (Kaufman and Rousseeuw 1990, 1). Mirkin makes the epistemological case for clustering as a knowledge-discovery process by noting that classification, under which clustering falls, helps us “shape and deepen knowledge, capture the structure of phenomena and relate different aspects of a phenomenon in question to each other”. He suggests that “clustering should be considered as classification based on empirical data in a situation when clear theoretical concepts and definitions are absent and the regularities are unknown,” and that the process to find and describe clusters validates classification as knowledge discovery (Mirkin 2005, 35-6). Cluster analysis falls under the umbrella of statistical data analysis, and has been used in fields as wide as social science, computer science, marketing research, and bioinformatics. Data for cluster analysis can come in two forms. There are *n* objects (in our case, accounts) that have *p* attributes, which give us an *n*-by-*p* matrix, where we can attempt to find groupings based on the similarities in attributes. The second is an *n*-by-*n* matrix where we attempt to find pairwise similarities or differences for each object, with every other object (Kaufman and Rousseeuw 1990, 4).

To quantify their degree of dissimilarity, we compute the “distance” between the objects. The two most popular distance measures for two objects *i* and *j* with *p* attributes on a continuous scale are the Euclidean distance:

$$\mathrm{d}(i, j) = \sqrt{(x_{i1} - x_{j1})^2 + (x_{i2} - x_{j2})^2 + \cdots + (x_{ip} - x_{jp})^2}$$

And the Manhattan distance:

$$\mathrm{d}(i, j) = |x_{i1} - x_{j1}| + |x_{i2} - x_{j2}| + \cdots + |x_{ip} - x_{jp}|$$

A generalization of both is the Minkowski distance, where the calculations above are extended to q dimensions, where q can be any real number equal to, or greater than 1:

$$\mathrm{d}(i, j) = (|x_{i1} - x_{j1}|^q + |x_{i2} - x_{j2}|^q + \cdots + |x_{ip} - x_{jp}|^q)^{1/q}$$

This is often called the L_q metric, with Manhattan and Euclidean distances being the L_1 and L_2 metrics respectively.

These distances in particular, and any dissimilarity metric in general, has the property that they "are small when *i* and *j* are 'near' to each other and that become large when *i* and *j* are very different." (Kaufman and Rousseeuw 1990, 11-13, 16)

Clustering algorithms come in two major flavors – partition-based and hierarchical. A partitioning method constructs *k* clusters, where each cluster contains one or more objects, but an object can belong to only one cluster. These algorithms generally try to find clusters such that objects in the same cluster are closer to each other than objects in different clusters (Kaufman and Rousseeuw 1990, 39).

Hierarchical clustering offers all possible values of *k* in the same run, in that the output consists of $k = 1$ where all objects are in the same cluster, and $k = n$, where all objects are their own cluster. It also yields clusters such that the difference between $k = r$ and $k = (r + 1)$ is that "one of the *r* clusters splits up in order to obtain $(r + 1)$ clusters." There are two approaches to creating hierarchical clusters – agglomerative, where all objects start as their own cluster and are then merged into fewer ones based on how not dissimilar they are, and divisive, where all clusters start as one, and are split until we end up with *n* clusters (Kaufman and Rousseeuw 1990, 44).

It may seem as though hierarchical clustering provides all possible *k* clusters, allowing us to explore all possible number of account groups in one go, but it has its drawbacks. While partitioning attempts to create the "best" clustering through repetition and can correct errors, hierarchical clustering algorithms can only offer rigid merges or splits, which ossifies erroneous decisions (Kaufman and Rousseeuw 1990, 44-45). Our review of literature does not suggest that there is a hierarchical order to savings behavior either, but that they are more in classes which cannot be easily subsumed into another, which conceptually aligns more with the partitioning model.

A partitioning algorithm generally goes as follows:

1. Given a set of objects and a dissimilarity measure between them, we start by selecting k representative objects in our initial k clusters. These starting "seeds" are often randomly assigned.
2. We then assign each object to a cluster based on minimized dissimilarity to the representative object.
3. We reassess the representative object once all objects have been assigned a cluster, and reassign the representative object as necessary to reduce the total amount of dissimilarity within a cluster.
4. If the total dissimilarity for all clusters is below a certain threshold, we consider the exercise complete. Otherwise, we repeat the process from step 2.

The two key determinants of this process are the dissimilarity measure, and the reassignment of the representative object. We use the Euclidean distance, or L_2, as our dissimilarity measure, mostly because there is no good reason to use a different one than what is usually considered to be the default measure for the distance between two objects in n-dimensional space.

There are essentially two common options for reassignment of the representative object. The k-means algorithm takes the centroid of the clusters assigned in step 2 above by averaging the measurement values along each of the dimensions. For a cluster v with n_v objects, the f^{th} coordinate of this new representative object is given by (Kaufman and Rousseeuw 1990, 112):

$$\bar{x}_{f,v} = \frac{1}{n_v} \sum_{i \in C_v} x_{i,f}$$

Note that this centroid does not necessarily have to be a member of the original dataset – it is any point in space that has the least average distance to all the objects in the cluster.

The k-median algorithm differs in that it calculates the dissimilarity measure, in our case, L_2, for each pairwise object in a cluster, and reassigns the representative object such that the average distance of that *object* to all the other objects in the same cluster is minimized. Thus, unlike *k*-means, *k*-median clustering always assigns a representative object that is a member of the original dataset. We choose to use the k-median algorithm as it is "more robust with respect to outliers" (Kaufman and Rousseeuw 1990, 40-41). During the actually clustering exercise, we ran both *k*-mean and *k*-median algorithms, and confirmed that *k*-medians provided smoother shapes for clusters (seen in Figure 12) compared to the jagged features evident in *k*-means outputs due to outliers.[1]

In terms of our dataset, we intend to cluster comparable subsets of balances. To do this, we will have to create comparable subsets and determine what an appropriate number of clusters would be. The following two sections do just that, before we turn to identifying clusters that give us our motifs.

Splicing Segments

Before we can start looking for patterns, we have to identify segments that contain each of these patterns. Based on our understanding of the purpose of savings accounts as being one to build

[1] What Stata calls as the k-median algorithm is also known as the k-medoid algorithm, since the dissimilarity-measure-minimized objects are known as medoids.

up a lump sum which is then drawn down, we consider a segment as being contained within that period of time from when the build-up of this lump sum starts, to when it terminates. Or quite simply, the period between subsequent zero balances.

Identifying these breaks where balances approach zero is not a trivial problem to solve, for a various reasons:

1. Many accounts will not reach zero balance often, if ever, but will get to relatively low balances compared to balances maintained at some other point in the same segment.

2. “Near-zero” is relative. $10 is “near-zero” for an account where balances reach $10,000, but not for an account where the balance only ever reaches $50.

3. The same account could have segments where the maximum balance is $100, and segments where the maximum balance is $1,000. Since we are looking at multiple years of data, we would want a concept of “non-zero” that is adapts to contextually relevant balances.

4. Some accounts will never reach a balance close to zero – there will be a significant residual balance over which funds are accumulated and drawn down.

We will use the term “near-zero balance” to mean “at or near zero balance” throughout this piece for the sake of brevity. What follows is a description of our attempts to identify these breaks.

The initial attempt at identifying segments built on the concept of

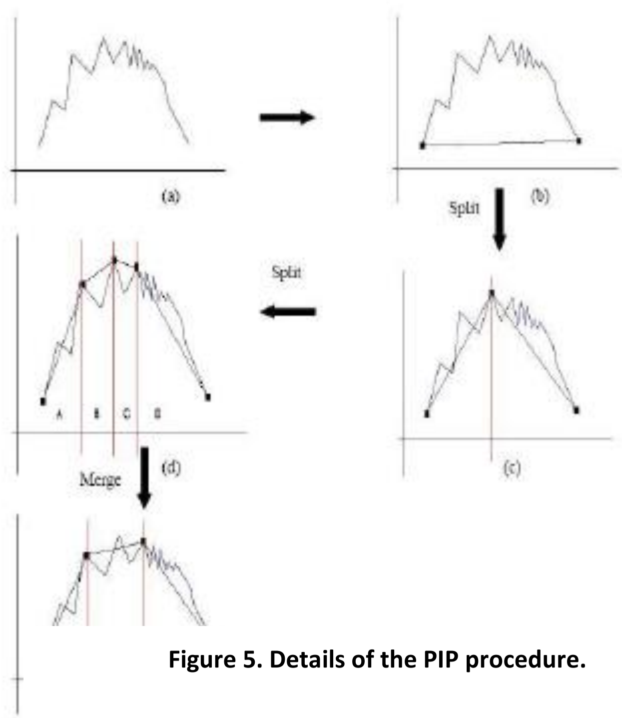

Figure 5. Details of the PIP procedure.

Perceptually Important Points (PIPs) that was first introduced by Chung et. al in a paper on time series data mining (Chung, et al. 2001). The PIPs approach was designed particularly with financial data in mind, and at first glance, seemed quite appropriate to our challenge. We attempted to adapt them to identify segments because PIPs are specifically designed to identify structural discontinuities.

The basics of the process are outlined in Figure 5, where the image used by Chung et. al. is reproduced exactly through screen capture (Chung, et al. 2001). The process initializes by starting at the edges, as in Figure 5 (b). It then finds the point along the time series which has the greatest Euclidean distance from the line connecting the two PIPs, as seen in Figure 5 (c). It proceeds by continuing to find PIPs between subsequent adjacent PIP pairs by calculating Euclidean distance, as seen in Figure 5 (d). A merge process cleans up the PIPs by removing adjacent sub-segments with similar slopes, as seen in Figure 5 (e).

Once the PIP algorithm was implemented in Python, checks showed that they were very good at finding peaks where balance was maximized, and less so at finding absolute troughs. Attempting to use the logarithm of the savings amounts helped somewhat, but did not ameliorate the situation. We also attempted to ignore all the peaks, by adapting the merge process to throw them out and stitch adjacent sub-segments together, but the results were, again, not acceptable. Finally, we attempted to find near-zero balances by identifying the Euclidean distance from not just the edges of the segment, but also from the maximum balance value. The use of the third distance calculation from the maximum represents a modification of the PIP algorithm. This worked well enough when the maxima was about equally spaced from the edges, but tended to return matches where cutoff was near one of the edges when the maxima was situated near the other edge. In the end, PIPs were substituted with a self-designed approach described below. Even though we did not use PIPs directly, we describe the framework as it provided the conceptual framework of successive identification of topographically salient points for our approach.

Our technique was quite simple but turned out to be surprisingly effective, in that is successfully identifies peaks and troughs in cases that would confound PIP. Within any given segment, initially the entire time series for a given account, it finds the maximum and minimum balances. As long as the minimum balance is below a certain threshold percentage, the segment is cleaved at that incidence of minimum balance to form two sub-segments. Each sub-segment is treated as a segment, the minimum-as-a-percentage-of-maximum check is conducted and that segment is cleaved in turn if it falls below the threshold. For example, if we decide on a threshold of 5% and have an account that has a maximum balance of $1,000, we will split up that segment for the lowest balance that is lower than $50. Now that the balance pattern is split into two, we look at each half separately and calculate the same again – if the minimum is less than 5% of that segment's maximum, the segment is cleaved. We continue until the minimum balance is no longer less than 5% of the maximum balance in that segment.

a. Account (a)

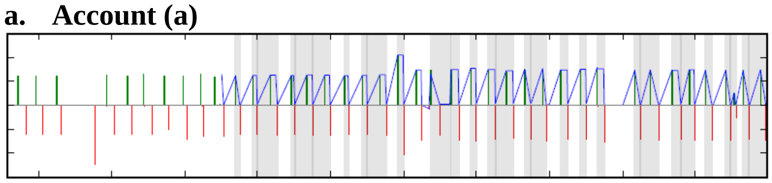

Number of segments identified: 32

b. Account (b)

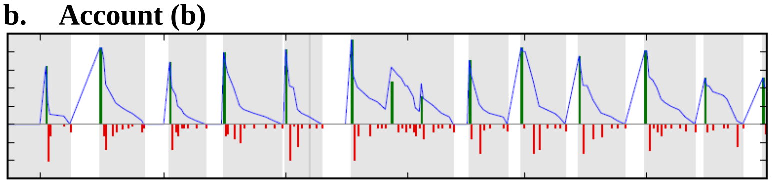

Number of segments identified: 13

c. Account (c)

Note that we don't split segments every time the 5% threshold is crossed, but only at the minimum, so that each segment creates two sub-segments at most. This is because it is possible that adjacent balance values could all be below the threshold, and splitting on all of them would lead to stunted, spurious segments. This threshold based splicing also allows for local maxima balances to stay relevant. Consider an account with balances of $1,000, $500 and $100, separated by periods where balances are in the low single digits. As long as $100 is separated from both $1,000 and $500 by a balance of $5 or less, it will show up on its own in a separate segment.

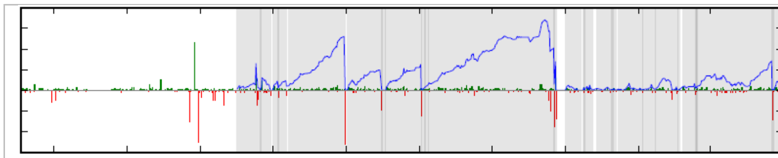

Number of segments identified: 34

d. Account (d)

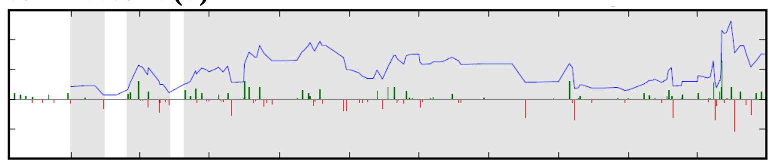

Number of segments identified: 3

Figure 6. Segments created by progressive cleaving at 10% threshold.

After trying various values for this threshold between 5% and 25%, we settled on a value of 10% for the threshold for this iteration. 10% seemed to provide a balance between a situation where dips were missed because they were just above a 5% threshold, and one where dips were seen but were not really there, for 25%. A 10% threshold creates a situation where a pattern has terminated once the balance falls 90% or more compared to the maximum balance seen before that point of time. The results of this algorithm splicing at a 10% threshold are presented in Figure 6**Error! Reference source not found.** for the four accounts presented in Figure 4 earlier. Each segment identified is labelled in grey.

Visual inspection suggests that the segment identification algorithm is working quite well. Each grey band represents a segment. There clearly is a plethora of segments for at least two of the four accounts – (a) and (c). The segments for (a) and (b) are also fairly regularly spaced, reflecting the underlying transactional regularity. (a) and (c) do not have grey bands denoting segments in the first third of their time series because the balance data from 2010 was not available for these accounts.

Note that the number of segments is a function of how active accounts are. Since we chose fairly active accounts for illustration purposes, it is quite possible that the typical account in the overall sample may have fewer segments.

Overall, 413,845 segments were identified from the 70,994 accounts in our sample. Accounts can have a range of segments, and segments can have a range of balance values. They can also span various lengths of time. Distributions of segments by account are presented in Table 3, those of balances by segment are presented in Table 4, and lengths of segments in days in Table 5 below.

Number of segments per a/c	**Number of accounts**	**%**	**Cumulative %**
1	35,414	50%	50%
2	7,680	11%	61%
3	4,446	6%	67%
4	3,196	5%	71%
5	2,327	3%	75%
6	1,815	3%	77%
7	1,466	2%	79%
8	1,310	2%	81%
9	1,080	2%	83%
10	918	1%	84%
More than 10	11,329	16%	100%
Total	70,981	100%	

Table 3. Number of Segments per Account

More than half the accounts have only one segment, implying that less than half the accounts deplete or almost deplete funds for the period of time we have data for. This implies that when we run clustering algorithms on segments, much of the comparisons will be across accounts, which in turn means that we must present segments to the clustering routing in a form that allows them to be compared across accounts.

Number of balance values per segment	Number of Segments	%	Cumulative %
1	10,651	3%	3%
2	123,464	34%	37%
3	72,313	20%	56%
4	44,576	12%	69%
5	28,212	8%	76%
6	18,792	5%	81%
7	13,109	4%	85%
8	9,456	3%	88%
9	7,193	2%	89%
10	5,616	2%	91%
More than 10	32,899	9%	100%
Total	366,281	100%	

Table 4. Number of Transactions per Segment

We have a wide variety of balance data points per segment. We can see this is possible from Table 4, where (a) mostly has only 2 balance values per segment, while (d) has more than 10 in its longest segment. This variability adds further considerations that need to be addressed to ensure that segments with different numbers of balance values can be compared to each other.

Length of Segment (months)	Number of Segments	%	Cumulative %
< 1	282,160	68%	68%
1	41,767	10%	78%

2	15,522	4%	82%
3	8,646	2%	84%
4	6,078	1%	86%
5	4,231	1%	87%
6	4,459	1%	88%
7	2,529	1%	88%
8	2,171	1%	89%
9	1,839	0%	89%
10	1,573	0%	90%
11	1,470	0%	90%
12	1,431	0%	90%
> 12	39,969	10%	100%
Total	413,845	100%	

Table 5. Length of Segments (Months)

More than half the segments are less than a month in duration, and three quarters are less than three months in duration. One out of ten segments extends for over a year. There seems to be some variation in terms of the length of the period that a segment manifests itself over.

Balance per Segment	10^{th} Pctile	25^{th} Pctile	Median	75^{th} Pctile	90^{th} Pctile
Maximum (KES)	300.00	1,946.77	6,301.43	16,379.96	42,648.40
Average (KES)	177.62	966.35	3,342.65	8,760.20	23,409.17
Maximum (USD)	3.33	21.63	70.02	182.00	473.87
Average (USD)	1.97	10.74	37.14	97.34	260.10

Table 6. Distribution of Segment Balances, Maximum and Average

Segments can have quite a range of balance values. Table 6 presents the range of average and maximum balance values within segments over various percentiles to illustrate the range over which they can exist. For example, we see above that the average balance for the middle 80% of accounts varies between USD 2 to about USD 260. An exchange rate of USD 1 = 90 KES is used as an approximation as the exchange rate has fluctuated between 80 and 100 over this time period.

Given the complexity of the process and the need to handle accounts with an arbitrary number of data points, we utilized the programming language called Python to generate these segments. The code utilizes a technique called recursion that ensures that the most salient cleavage points are handled first, and then progressively moves to less significant ones until some tolerance threshold is reached.

Harmonizing Balances

Now that we have our segments, the next step requires finding if the areas within the grey bands in Figure 6 have recurring patterns, within the same account and across accounts. Before we can do that, however, we need to make sure that they are in a form where they can be compared to each other. There are three issues that prevent us from doing so immediately:

1. The number of times balances change within a segment can vary from one to a few score. We saw this in Table 4 above.

2. The time period over which a segment manifests itself can be as little as a few days and as much as many months, with an arbitrary number of changes to balances in that period. We saw this in Table 3 above.

3. The maximum balances of these accounts can vary wildly, over orders of magnitude. This is illustrated in Table 6 above.

We address (1) and (2) by massaging the arbitrary number of balance data-points into one with a fixed-length through some manner of extrapolation. We address (3) by converting absolute balances to percentages, where the values are percentages of the segment's maximum balance. Regular mean normalization is not attempted because we don't know what the underlying distribution of these balance patterns are (in fact, if we did, that would be the answer to the typology question!). As with splicing segments, this harmonization task was also carried out utilizing Python code we authored.

Let us look at each of the two harmonizing steps in detail.

Step 1 – Fixed Number of Balance Values for All Segments

First, we consider the conversion of balance sequences of arbitrary lengths into one of specific length that addresses concerns (1) and (2) above. Every segment, irrespective of the number of balance values it has, and the number of days it spans, will be converted to a segment of fixed size n. If every segment is of length n, a pairwise comparison of each value becomes a much simpler exercise compared to uneven segments. It is also necessary for objects to have the same dimensions when dissimilarity/distance measures are calculated as part of the clustering process.

Irrespective of whatever value we choose for n, we will need to split up sub-segments for segments with less than n balance values, and merge sub-segments for those with more than n balance values. We would ideally like to reduce transformations as much as possible to preserve original content. Let us illustrate this process below, with $n = 5$.

A simple approach of achieving harmonized sub-segment counts could have been to divide the time series up into five parts, each of equal length, and then finding the weighted average balance in each segment, as is sometimes suggested (Tanaka, Iwamoto and Uehara 2005). This has the unfortunate effect of creating patterns where they don't exist. Let us consider a hypothetical situation where an account has $100 in the account for the first 55 days, and $1 for the last 45 days. If we wanted to break this account up into 5 equal parts, the first two sub-segments will have a balance of $100, the last two sub-segments will have a balance of $1, and the sub-segment in the middle will have a balance of ($100 * 15 + $1 * 5) / (15 + 5), or $75.25.

What was a simple step function has now become a staggered balance pattern with three distinct steps, as shown in Figure 7 below.

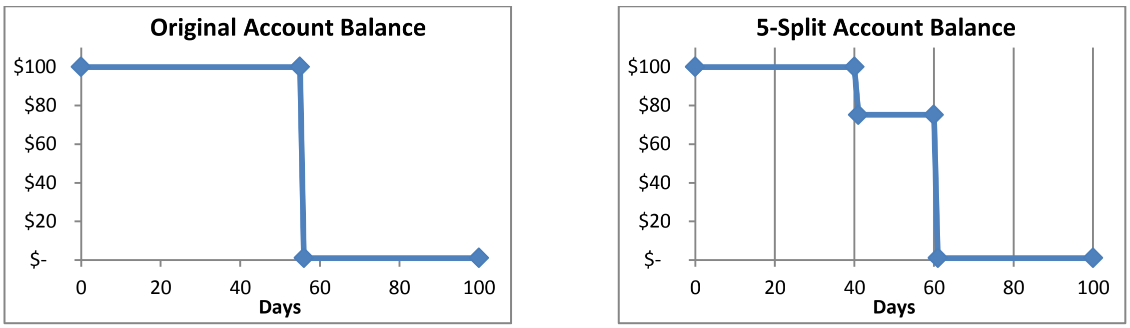

Figure 7. Splitting Balances into Equal sub-segments

We must endeavor to preserve the balance profile "shape" as best as we can while we try to arrive at balance sequences of the same length, and prevent as much loss by adjustment as possible to prevent generation of spurious balance values. One key to that is not letting additional artifacts to appear.

The algorithm we implement dispenses with the notion that each of the five sub-segments will be of the same length. Instead, it tries to make them "as equal as possible". The approach is different depending on if the segment has more than, or less than five balance values. If it has exactly five balance values no transformation is applied in this step.

If we have **less than** five balance values, the longest segment for which a balance exists is identified, and cleaved into two, with each half being assigned the same balance. If the new balance segment is still less than five-values long, the next longest segment is identified, cleaved in half, and assigned that balance value. And this cleavage continues until we arrive as exactly five balance values.

This is illustrated in Figure 8 below for one of the accounts in our sample, where an account has a balance of 16,290.30 for 33 days, 32,641.30 for the next 22 days, and ends with a balance of 45.30.

Days	Balance
33	16,290.30
22	32,641.30
End	45.30

⇒

Days	Balance
16.5	16,290.30
16.5	16,290.30
22	32,641.30

⇒

Days	Balance
16.5	16,290.30
16.5	16,290.30
11.0	32,641.30
11.0	32,641.30
End	45.30

Figure 8 progressive cleavage of balance segments to arrive at 5 balance values.

Note that we consider periods consisting of fractional days acceptable, partly out of necessity as converting a five balance sub-segment that occur over less than five days must necessarily require some of the days to be less than one day long. Also, note that the final balance is a fixture that does not have a day value – the pattern terminates when this value is reached.

If we have **more than** five balance values, the balances that exist for the shortest amount of time are merged into adjacent segments. Thus, the minimum length balance sub-segment is identified, and merged to the shorter of its two adjacent sub-segments. Then, the next shortest sub-segment is found, and merged to its shorter of two adjacent sub-segments, and so on, until we have built up to just five balance values. This merging process is illustrated in Figure 9 below for an account in our sample.

Days	Balance
12	7,200
12	6,675
32	4,645
5	11,245
22	10,220
8	5,650
87	2,105
End	1,075

⇒

Days	Balance
12	7,200
12	6,675
32	4,645
27	10,410
8	5,650
87	2,105
End	1,075

⇒

Days	Balance
12	7,200
12	6,675
32	4,645
35	9,321.86
87	2,105
End	1,075

⇒

Days	Balance
24	6,937.50
32	4,645
35	9,321.86
87	2,105
End	1,075

Figure 9. Progressive Merge of Balance Segments Down To Five Balance Values.

Step 2 – Balances as Percentage of Maximum Balance

Now that we have a process to arrive at exactly five balance values, we can harmonize their amplitudes to arrive at a percentage based representation by dividing all values by the maximum in the sequence. This is a fairly trivial process, and is illustrated in for the two accounts presented above in Figure 8 and Figure 9.

Balances from progressive cleavage

Days	Balance	Percentage
16.5	16,290.30	0.499
16.5	16,290.30	0.499
11.0	32,641.30	1.000
11.0	32,641.30	1.000
End	45.30	0.001
Maximum	32,641.30	

Balances from progressive merging

Days	Balance	Percentage
24	6,937.50	0.744
32	4,645.00	0.498
35	9,321.86	1.000
87	2,105.00	0.226
End	1,075.00	0.115
Maximum	9,321.86	

Figure 10. Harmonizing Balance Values to Percentages

These percentage values are the input that will be used to find patterns. Note that this manner of flattening balances that could be orders of magnitudes apart is not completely without cost as it prevents us from taking the balance values itself into consideration as part of the typology creation exercise.

For this thesis, we choose ***n* = 6**.

The primary motivator for this derives from the fact that reducing the number of segments as part of the transformation leads to some loss of information, as we will see later in this section. About three-quarters of the segments have less than six sub-segments, requiring them to sprout additional ones, and a little less than 20% will need to merge some of their sub-segments into fewer ones (Table 4). One could arguably suggest that since a higher number of *n* leads to less merging of sub-segments and therefore loss of information through smoothing, we should take a much higher number that encompasses a greater proportion of segment. We have explored a value of *n* up to 10, and do not find significant additional benefit in adding additional sub-segments where the same value is replicated multiple times across spliced segments.

Of the 413,845 segments we identified for 70,994 accounts in the previous section Splicing Segments, 355,574 segments for 41,365 accounts are successfully harmonized.

Of the 58,271 segments which were not harmonized, 29,923 segments belonged to 29,923 accounts for which we do not have transaction data – they are dormant, and have a static value for the duration of the segment which also spans the duration for which we have balance data for. The remaining 28,348 segments consist of only one balance data point throughout the segment. Both cases represent segments where there is no activity. The difference between the two is that in the latter groups, the account has other segments where there *is* activity, and can be utilized for our purposes. These 58,271 segments are ignored during the harmonization process, and subsequent motif creation.

It should be noted that not being to include these dormant segments has the effect in setting up a truncated regression model, as it is missing observations for all accounts that have conducted zero transactions during the two and half years (Wise and Hausman 1977). This implies that our technique does not take into account the impact of dormant accounts, and cannot be used to inform us about accounts with zero activity. We find this constraint acceptable as we could not construct segments if there were no changes in balance levels anyway, and point this out more to note that this technique is applicable only if there is at least some account activity.

Identifying Recurring Patterns (Motifs)

The number of clusters would emerge from the harmonized segments is not immediately obvious. We rely on two criteria to arrive at the number of clusters we select, and therefore how many motifs we identify. Quantitatively, we rely on the Calinski–Harabasz pseudo-F index to guide us to an optimal number of clusters, since “a larger value indicates more distinct clustering” (StataCorp 2009). Thus, we have to identify the number of clusters for which this pseudo-F index score is highest. Qualitatively, we look at the cluster means and see if they are consistent with we would expect from our literature review and familiarity with anecdotal behavioral patterns. Let us take each in turn, as we consider between three and ten clusters. The upper bound may seem somewhat arbitrary, but as we shall see below, it soon becomes evident that a higher number of clusters do not lead to more distinct clustering.

There is one additional complication that prevents us from simply running k-median cluster analysis for two to ten clusters, and comparing their scores. Unlike hierarchical clustering, the outcome of the k-medians is not deterministic, i.e. we are not guaranteed the exact same clusters every time the same algorithm is run on the same dataset. This is because the final clusters are dependent on the initial cluster seeding, which in itself is random. It is possible that some segments are assigned to different clusters, based on a different starting point. We therefore run the clustering algorithm thirty times, for each cluster count between three and ten, taking the pseudo-F score from each of the runs, and comparing them in aggregate.

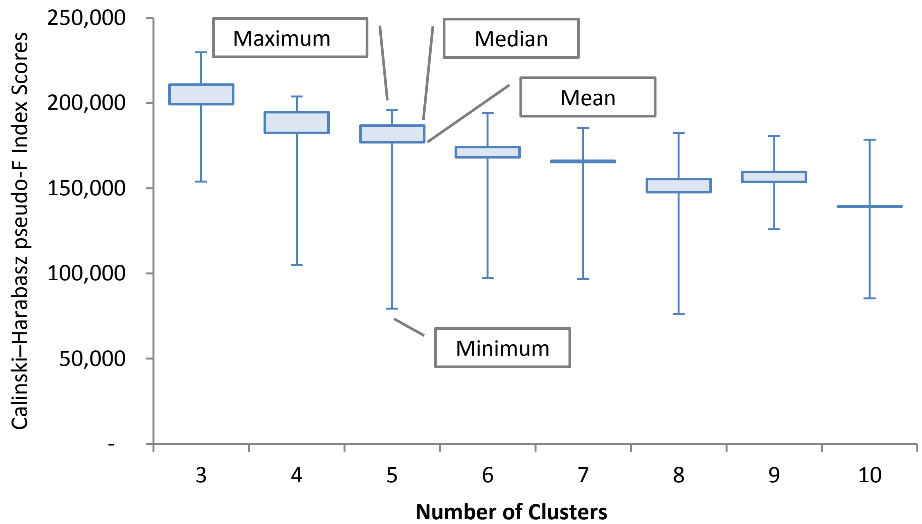

Figure 11. Calinski–Harabasz pseudo-F Index Scores, by Cluster

We can inspect the mean, median, minimum and maximum values for the pseudo-F scores for each of the cluster sizes between three and ten in Figure 11 (all scores are available in detail in *Appendix A. Calinski–Harabasz pseudo-F Index Scores*). The scores do not mean anything per se, and comparisons between scores are simply one of order. The maximum, mean and median scores monotonically decrease from three clusters to eight. Nine clusters edges slightly higher before scores fall for ten clusters. The pseudo-F scores are therefore telling us that the lesser the number of clusters, the more distinct the clustering.

Let us see if a qualitative inspection of what these clusters actually look like support this notion that we should take the smallest number of clusters. We calculate the average percentage value for each of the five sub-segments, for each of the cluster counts with the maximum pseudo-F scores. This gives us a visual indication of what shapes each of the clusters are centered around, presented in Figure 12 below.

a. Three Clusters

b. Four Clusters

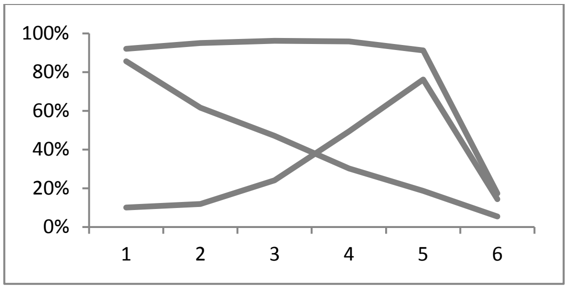

Pseudo-F score: 143,988

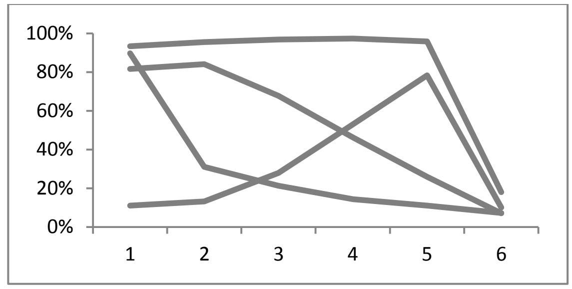

Pseudo-F score: 140,135

c. Five Clusters

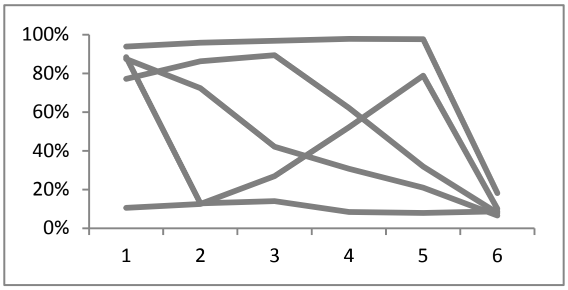

Pseudo-F score: 135,151

d. Six Clusters

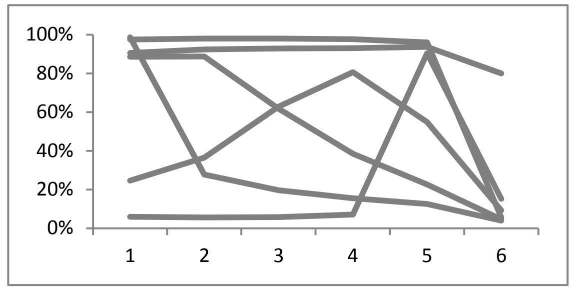

Pseudo-F score: 132,760

e. Seven Clusters

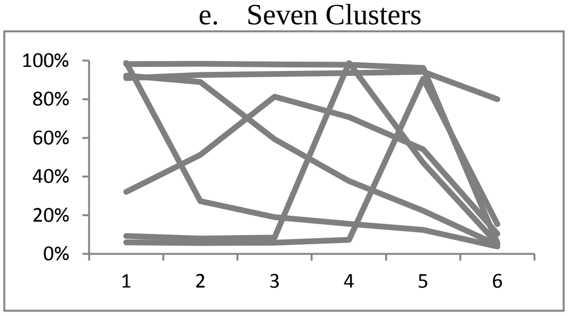

Pseudo-F score: 135,697

f. Eight Clusters

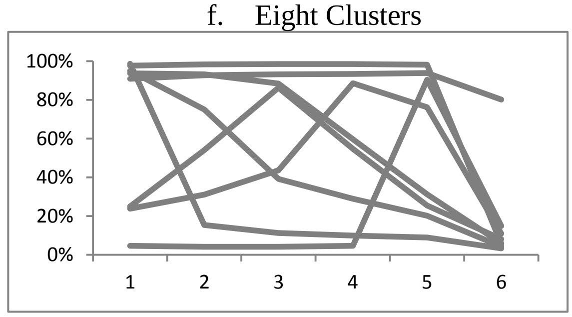

Pseudo-F score: 133,491

g. Nine Clusters

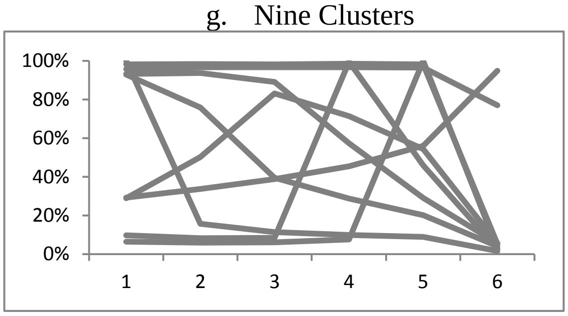

Pseudo-F score: 134,316

h. Ten Clusters

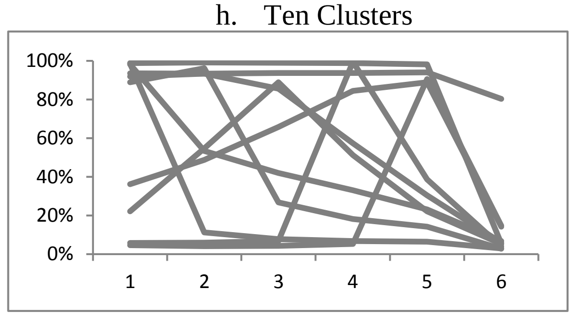

Pseudo-F score: 136,704

Figure 12. Pattern Clusters and their Calinski–Harabasz pseudo-F Index Scores.

It is interesting to note the evolution of cluster formation. Comparing (a) and (b), we see the pattern with the uniform gradient that is decreasing in value from left to right for three clusters gets split into two patterns, one with a higher initial gradient, and a lower initial gradient. In (c), we see the line in (a) reappear for five clusters. Additional clusters beyond three continue to provide distinctive patterns, so we cannot settle for three clusters only, despite their highest pseudo-F scores.

From (d) onwards, however, it seems that the patterns become less distinctive. Additional clusters identified shadow existing clusters and in some cases are simply phase shifted, where the peak appears a sub-segment or two earlier, but otherwise have the same shape. By nine clusters, it is difficult to tell what is distinctive any more. Settling for five clusters therefore appears to be a good compromise of letting pseudo-F scores drop as a price for identifying patterns with diminishing distinctive features.

Thus, we settle for five recurring patterns, or “motifs,” henceforth.

Let us take a look at each of the five clusters in more detail in Figure 13 below. We are simply taking the information presented in Figure 12 for the case with five clusters, and drawing a bar chart for each of the colored lines so that we can visually inspect at each motif separately. We also include an “average” pattern, which represents the percentage values of all the segments, for each of the five sub-segments for reference. The numbering of the motifs is arbitrary.

The most prevalent motif is “Motif 2”, consisting of a cluster formed off of 43% of the harmonized patterns. This seems to represents account balances where most of funds are pulled out in one withdrawal. “Motif 5” is the next most prevalent, at 20%, where a withdrawal seems to take place at a Sustained pace. “Motif 3” is the least prevalent, where most of the funds are taken out initially, with some residual. We will look at these motifs more closely in the next section, and assign more intuitive names.

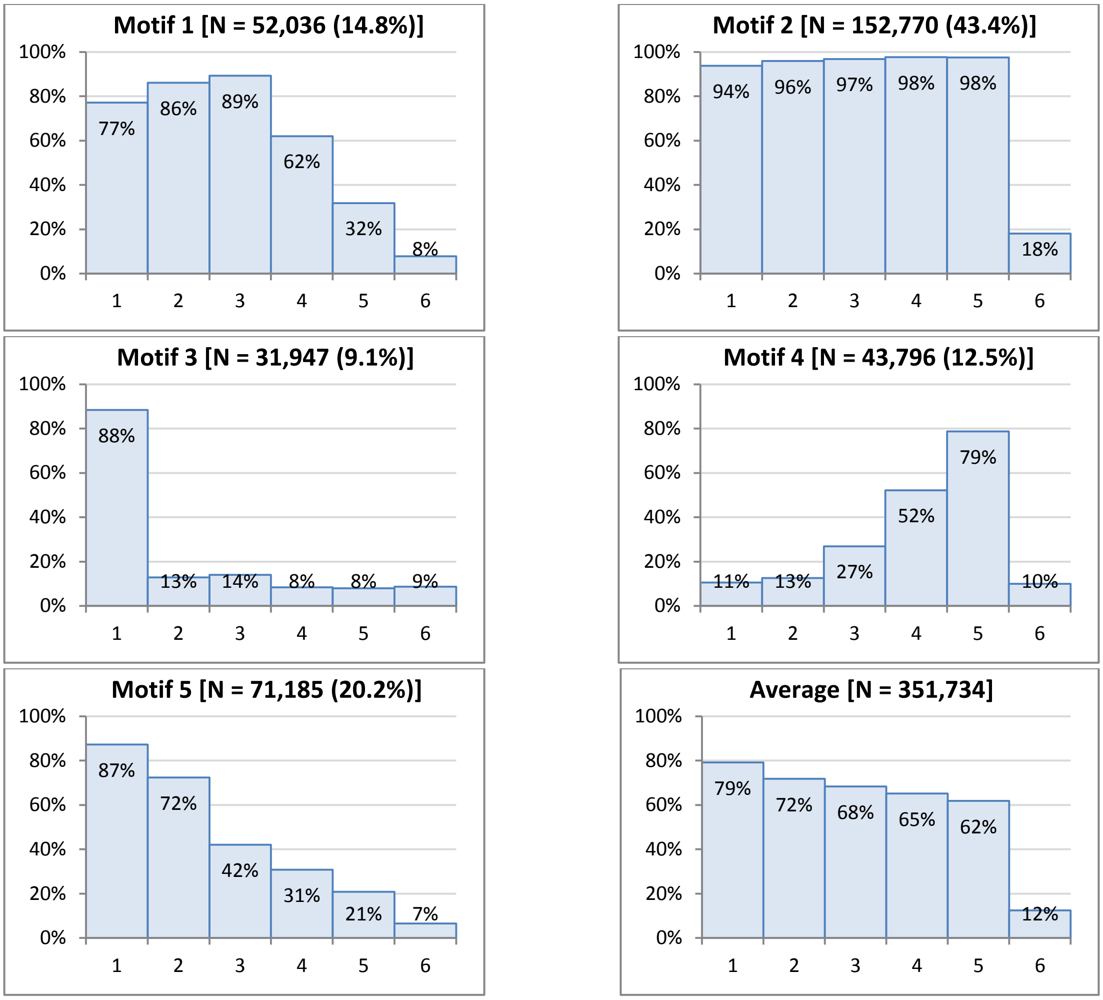

Figure 13. Details of Each of Six Clusters/Motifs.

Assigning Motifs to Segments

Technically, motifs have already been assigned to segments by Stata. Every segment for every account was used to come with the clusters, which means each was also assigned to a cluster. Upon visual inspection, it was evident that some small, but non-trivial number of segments was being assigned clusters where they really did not belong – visually, their profile did not fit the motif.

To account for this, we carry out an additional step where we find the best motif match to the segment at hand using a Euclidean distance measure. Recall that the Euclidean distance, d_e, is simply the square root of the sum of squared differences:

$$d_{e,m} = \sqrt{\sum_{i=1}^{6}(y_i - x_{m,i})^2} \qquad \text{(ii)}$$

Here, *y* represents the harmonized segment we are attempting to match to a cluster, while x_m represents each of the *m* clusters/motifs. The summation goes from one to six as there are six sub-segments.

We apply one additional filter before assigning motifs. Using the entire dataset of segments, we calculated what the "average" pattern looks like by looking at the percentage-wise balance levels at each of the six sub-parts of a segment (Figure 13). When motif matching is carried out, distances to this "average" pattern is carried out too, and a segment is assigned a motif *only if* it is closer to the motif than it is to the average pattern. This turns out to be surprisingly effective at weeding out balance patterns that were assigned clusters by Stata, but whose membership of said cluster was visually suspect. This exercise is carried out using Python code we authored.

Motif	No. of Segs	%
1	52,036	15%
2	152,770	43%
3	31,947	9%
4	43,796	12%
5	71,185	20%
Total	351,734	

(a) Initial Assignment by Stata

Motif	No. of Segs	%
1	44,145	13%
2	141,218	40%
3	32,689	9%
4	44,944	13%
5	55,836	16%
Unassigned	32,902	9%
Total	351,734	

(b) Modified Assignment

Table 7. Modified Assignment of Motifs

The results of this modified assignment to clusters are presented in Table 7. 9% of the segments that were assigned a cluster by the k-medians algorithm Stata are not assigned one of the five motifs as a result of this additional check. While impact is not completely uniformly distributed across the five motifs –with Motif 5s go down most, by 4%, and Motif 4s actually gain 0.3% of segments – it seems to impact all the motifs to some extent, suggesting that this "correction" is not punitive to any particular motif alone.

The motifs represented in Table 7 (b) represent the typology that we were seeking, and therefore are the end product of this pattern discovery process.

Figure 14 below provides examples of what this motif assignment "looks" like, helping us to intuitively grasp what is going on. Instead of calling them by non-intuitive numbers, we also assign them names. The names have been chosen with intent not to imply motive on the part of the client to any of the motifs. It also includes an example of what a rejected pattern looks like, though understandably rejected patterns will come in all shapes and sizes that are not close enough to one of the five motifs.

<table>
<tr><th>Stylized Motif</th><th>Matched Segment</th><th>Names</th></tr>
<tr><td>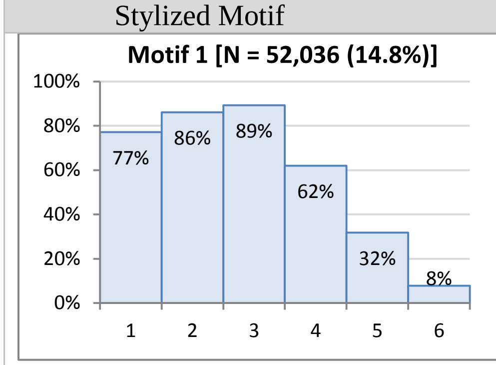
</td><td>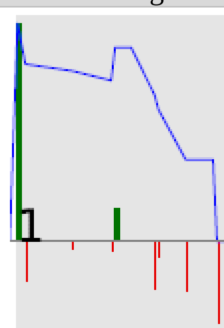
</td><td>Sustained Balance: A certain amount of balance is maintained for most of the duration of the segment, before it is drawn down. Some deposits may follow initial top-up. There are intermittent, small withdrawals.</td></tr>
<tr><td>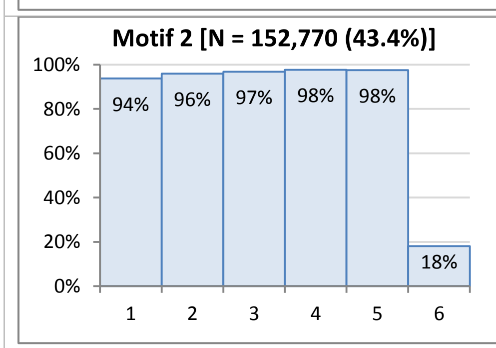
</td><td>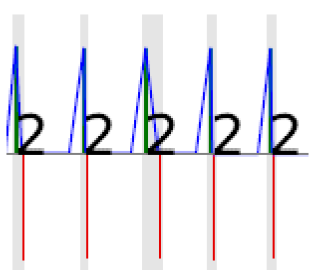
</td><td>Dump-and-pull: The amount deposited is pulled down in its entirety, or almost in its entirety. Withdrawal often occurs quite soon after deposit.</td></tr>
</table>

<table>
<tr><th>Stylized Motif</th><th>Matched Segment</th><th>Names</th></tr>
<tr><td>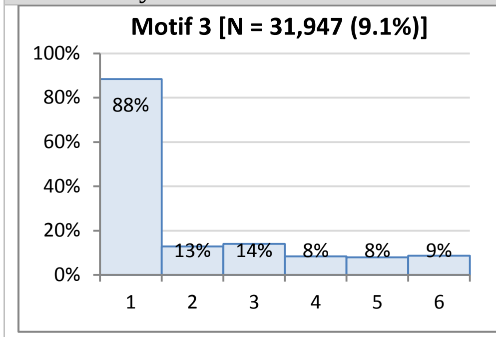
</td><td>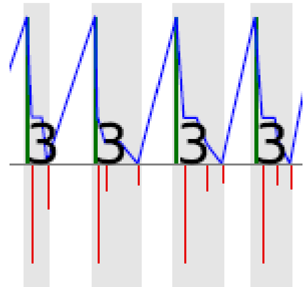</td><td>Fast Drawdown: A significant chunk of amounts deposited is withdrawn at the first opportunity, with the remainder of the balance slowly depleted thereafter.</td></tr>
<tr><td>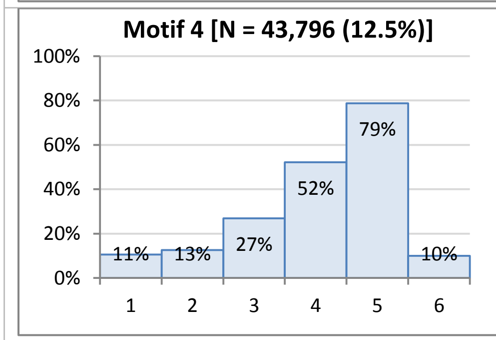
</td><td>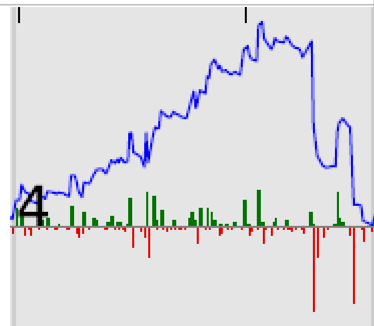</td><td>Accumulator: Small amounts of savings are slowly saved up to a certain lump-sum until they are finally drawn down. There may be intermittent, small withdrawals interspersed with deposits, but the general trajectory is of increased deposits.</td></tr>
<tr><td>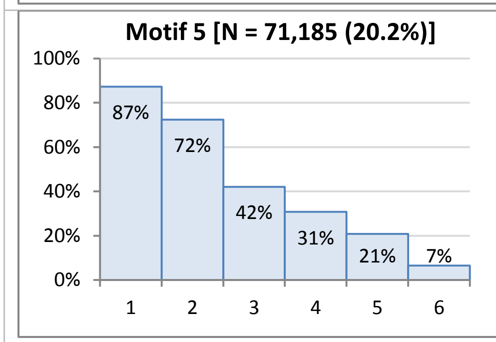
</td><td>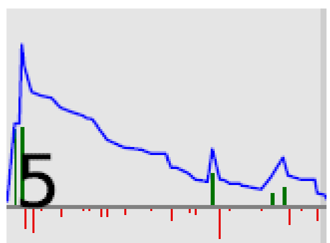</td><td>Slow Drawdown: This seems to be an intermediate profile between Fast Drawdown and Sustained Balance – initial deposit is withdrawn at a steady rate.</td></tr>
<tr><td>No Match (labelled as -1 by our code) – this happens when the segment fails the check of being closer to the Average profile than any of the five motifs.</td><td>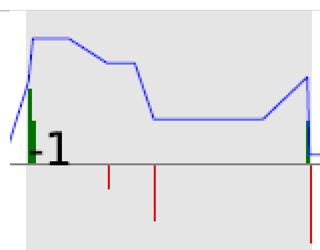</td><td>-</td></tr>
</table>

Figure 14. Segments Matched to Motifs

Given a long tail for these metrics (Figure 24), we look at the median values of average balance, number of deposits, number of withdrawals, average amount of deposits, average amount of withdrawals, and days covered in the segment (Figure 15). Note that the averages are calculated within segments. The medians are taken of those values *across* motifs manifesting themselves in those segments.

M1: Sustained Balance		**M4: Accumulator**	
Num. Deposits:	2	Num. Deposits:	2
Num. Withdrawals:	4	Num. Withdrawals:	2
Avg. Dep. Amt.:	KES 9,250 ($102.78)	Avg. Dep. Amt.:	KES 5,367 ($59.63)
Avg. W/d. Amt.:	KES 4,000 ($44.44)	Avg. W/d. Amt.:	KES 4,560 ($50.66)
Avg. Balance:	KES 7,982 ($88.69)	Avg. Balance:	KES 2,019 ($22.44)
Days Spanned:	22	Days Spanned:	29
M2: Dump-and-Pull		**M5: Slow Drawdown**	
Num. Deposits:	1	Num. Deposits:	1
Num. Withdrawals:	1	Num. Withdrawals:	5
Avg. Dep. Amt.:	KES 3,000 ($33.33)	Avg. Dep. Amt.:	KES 12,865 ($142.94)
Avg. W/d. Amt.:	KES 3,000 ($33.33)	Avg. W/d. Amt.:	KES 3,760 ($41.78)
Avg. Balance:	KES 3,112 ($34.58)	Avg. Balance:	KES 6,176 ($68.62)
Days Spanned:	4	Days Spanned:	20
M3: Fast Drawdown		**No Motif**	
Num. Deposits:	1	Num. Deposits:	3
Num. Withdrawals:	3	Num. Withdrawals:	4
Avg. Dep. Amt.:	KES 9,000 ($100.00)	Avg. Dep. Amt.:	KES 5,875 ($65.28)
Avg. W/d. Amt.:	KES 3,750 ($41.67)	Avg. W/d. Amt.:	KES 3,750 ($41.67)
Avg. Balance:	KES 1,619 ($17.99)	Avg. Balance:	KES 7,184 ($79.82)
Days Spanned:	19	Days Spanned:	33

Figure 15. Transaction and Balance Figures for Motifs

Motif 2, Dump-and-Pull, may not visually seem to fit the balances they are mapped to. Given how the algorithm works, the first leg representing the deposit is the one that is stretched, with the low threshold being pegged in the last sub-segment. This, the elevated balance level appears for the first five segments, followed by the withdrawal in the sixth segment. Importantly, this does not even look like "savings" behavior as there is no retention of funds, with a single deposit and withdrawal happening within 4 days of each other, on average. This suggests that the largest fraction of accounts simply using these no-frills accounts as current/transactional accounts.

Motif 4, Accumulators, are what we would understand to be the accretive savers – putting away small amounts of funds deliberately and over time to achieve a lump sum, which is then drawn down. There may be some small withdrawals along the way, but the general trend for funds accumulation is upwards. There accounts are what Rutherford had in mind in his "saving up" behavior. These motifs span the longest amount of time, at 29 days. About one in eight segments displayed this behavior.

Motifs 1, 3 and 5, what we are calling Sustained Balance, Fast Drawdown, and Slow Drawdown respectively, represent the “saving down” behavior. At a cursory glance, it may seem that the only difference between them is the rate at which the initial deposit is drawn down. It will become clear through the rest of the thesis that each of these motifs show somewhat distinct behavioral patterns and respond to changes differently. For now, we note some discerning features evident from the shape of their profile.

For Fast Drawdowns, their ability to save is somewhat limited, having needed to draw down most of the initial deposit in the very first withdrawal that soon follows said deposit. That they persist to save the remainder in this account may indicate perceived value of the account and/or desire to save from whatever limited residual funds these account holders have.

Slow Drawdown accounts seem to be funding regular expenses from an initial deposit of funds. Even though the withdrawals are regular, the amounts are not necessarily so, suggesting a comfort level with using the account as a funds repository with satisfactory liquidity. There are five withdrawals for every deposit for these motifs, and each lasts about three days.

Sustained Balance accounts also seem to display a certain amount of comfort keepings funds in the account, with occasional deposits and withdrawals that do not significantly change the balance maintained, till the need arises to draw it all down. They also have the highest median balance at almost KES 8,000, and manifest themselves over about three weeks.

Thus, it seems that Fast Drawdowns save what they can, when they can, Slow Drawdowns finance regular expenses using their savings account, and Sustained Balance ones use it as a funds repository.

The segments that could not be assigned to a motif as they were not close enough in shape to any of the five may be undefined in their profile, but it is worth nothing that they do display a tendency to accumulate funds, and effectively save. The average "no motif" account lasts longer than Accumulators, at 33 days, and has an average balance that is only second to Sustained balance accounts. These probably represent balance managers who intermediate funds while maintaining a certain level within the account.

Since we will be looking at balance, deposits and withdrawals, we present the distributions of average balance of accounts by deposit and withdrawal frequencies to see if it can provide us with any additional insight before we begin rigorous time series analysis. In the heatmap in Figure 16, median balance values are provided for the corresponding deposit and withdrawal frequencies. Overall, higher deposits and withdrawal levels seem to be correlated with higher balances, though necessarily in concert. We restrict frequencies to ten to prevent a gigantic table as this captures the vast majority of accounts.

ALL	**Withdrawals**									
Deposits	0	1	2	3	4	5	6	7	8	9
0	88	390	1,017	2,314	2,853	5,813	4,743	7,651	6,705	14,227
1	516	4,310	3,037	3,053	4,296	5,947	6,530	8,184	8,629	10,934
2	1,400	3,991	1,312	2,602	3,464	5,484	6,701	8,518	10,108	12,026
3	2,850	4,436	2,976	2,467	3,301	5,253	6,135	8,823	8,936	10,716
4	3,000	5,651	4,540	4,849	3,348	5,178	5,749	7,622	8,400	10,215
5	4,637	5,459	5,989	6,036	4,957	5,268	5,354	7,038	8,491	10,961
6	7,570	7,394	6,174	7,477	6,310	7,548	5,093	8,207	9,802	9,827
7	8,648	10,311	8,438	8,142	8,155	9,484	10,065	6,810	8,953	11,243
8	18,234	11,209	6,940	6,900	8,614	12,710	7,916	9,486	6,346	8,648
9	14,536	14,042	9,778	8,606	10,901	8,663	8,237	11,834	10,131	11,526

Figure 16. Avg Balance by Deposit and Withdrawal Frequency

Now, we break this down by motif and present the average balance heatmap by deposit and withdrawal frequency in Figure 17. Only buckets with more than fifty data points are shown, to prevent outliers from obfuscating overall patterns.

SB (1)	Withdrawals									
Deposits	0	1	2	3	4	5	6	7	8	9
0		2,445	3,533	4,681	6,136	8,679	5,202	11,461	7,369	
1	3,021	3,084	5,588	5,868	6,530	8,889	8,651	10,140	11,739	12,148
2		3,294	4,247	5,308	5,506	7,344	7,889	9,632	10,519	13,617
3		2,928	3,430	4,603	6,148	6,697	6,916	9,160	9,335	11,586
4			4,764	6,520	5,282	7,427	7,146	8,243	10,412	10,034
5			4,999	5,579	6,266	6,612	7,182	9,268	9,464	12,907
6			7,998	4,165	6,228	8,361	8,125	11,482	11,946	10,555
7				6,959		9,484	7,602		9,895	
8										
9										

DP (2)	Withdrawals									
Deposits	0	1	2	3	4	5	6	7	8	9
0	82	361	977	1,964	1,465	4,357	1,195			
1	515	4,795	3,105	3,189	4,502	9,075	7,483	14,110	20,688	
2	3,120	4,909	1,716	2,700	3,471	7,500	8,651	12,648		
3	7,994	6,460	4,628	2,098	3,145	5,537	9,362	18,245		
4	4,025	8,516	7,354	7,423	2,708	9,870	6,588			
5	5,711	10,607	6,786	12,679	13,852					
6		17,386	11,421				3,183			
7										
8										
9										

FD (3)	Withdrawals									
Deposits	0	1	2	3	4	5	6	7	8	9
0	-	753	1,312	1,620	2,407	3,752				
1	164	780	1,714	1,310	2,184	3,209	4,407	5,790	5,922	7,946
2	269	545	503	1,263	1,475	2,721	3,030	7,643	6,267	
3			1,122	582	1,507	4,017	3,379			
4					235					
5										
6										
7										
8										
9										

Acc (4)	Withdrawals									
Deposits	0	1	2	3	4	5	6	7	8	9
0	3,595	1,643	880	1,425	1,349					
1	481	857	968	1,106	1,128	1,385	1,503	2,337	3,146	
2	594	1,649	733	974	1,023	1,201	2,320	3,120	4,434	6,998
3	1,691	2,689	2,120	952	1,176	2,151	2,357	3,225	5,656	
4	1,947	4,509	3,831	2,989	1,576	1,289	2,258	4,520	7,274	
5	3,455	4,933	5,632	4,614	3,557	2,464	1,287		5,869	
6	6,049	6,770	5,004	8,807	7,009	5,146	4,079	1,134		
7	10,136	8,495	6,973	8,635	6,528		7,789			
8		9,600	7,954	6,886						
9		13,957		7,612						

SD (5)	Withdrawals									
Deposits	0	1	2	3	4	5	6	7	8	9
0			3,779	3,816	4,866	6,010	6,436	6,913	8,486	17,278
1		1,693	3,588	4,131	5,052	6,191	6,696	8,197	8,262	10,930
2		1,527	2,337	3,328	4,035	5,595	6,938	8,241	10,066	10,938
3			2,234	3,187	3,277	4,944	7,394	8,039	8,761	12,340
4			3,375	3,789	3,956	4,962	6,078	9,098	8,284	10,545
5					4,174	3,723	7,206	6,399	8,435	
6										
7										
8										
9										

Figure 17. Average Balance by Deposit and Withdrawal Frequency, by Motif

The heatmaps are generally consistent with the patterns the motifs purport to represent. For Accumulators, the balances are generally higher when deposit frequencies exceed those of withdrawals. For Sustained Balance, and Fast and Slow Drawdowns, higher withdrawal frequencies are associated with higher balances. This is possibly because higher amounts of initial deposits allow for a greater number of spaced out withdrawals. For all motifs, it is interesting to note that the lowest balances occur when the number of deposits and withdrawals are near each other.

Assigning Segments to Accounts

We complete the cycle of motif identification by assigning these motifs to each account. This process is illustrated below for the four accounts we have been using to showcase the various steps earlier in Figure 4 and Figure 6. The deposits (green), withdrawals (red) and balances (blue) are shown for the entire thirty months we have data for.

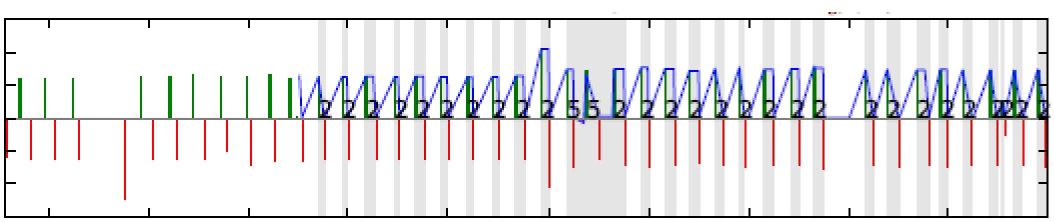

Slow Drawdown – 2 Dump-and-pull – 29

Figure 18. Motifs Assigned for Account (a).

Account (a) overwhelmingly engages in Dump-and-pull behavior, with 27 of the 29 segments (grey areas) being identified as such, but there are two segments where Slow Drawdown is taking place (Figure 18). The initial year does not have transaction data, and therefore no segments could be created.

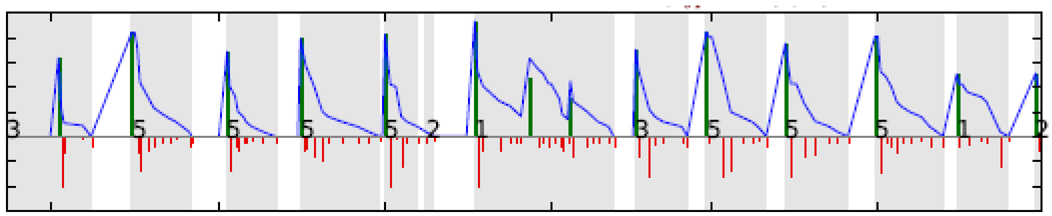

Sustained Balance – 2 Dump-and-pull – 2
Fast Drawdown – 1 Slow Drawdown – 5

Figure 19. Motifs Assigned for Account (b).

Account (b) is a mixed bag (Figure 19). There are five Slow Drawdowns, two Sustained Balances, two Dump-and-Pulls, and one Fast Drawdown. It is interesting to note that on visual inspection, it seems that this account is doing essentially the same thing, depositing a lump-sum and then drawing it down over multiple withdrawals. The timing and amount of the withdrawals make the difference, creating distinctions between Slow Drawdowns and the rest. Note that the Sustained Balance segment in the middle (labelled '1') consists of three deposits with interspersed withdrawals, and gets classified as such because the balance never gets low enough for multiple segments at the trough. Thus, the account seems to be engaging in the same behavior visually, but the motifs reveal nuances that suggest otherwise.

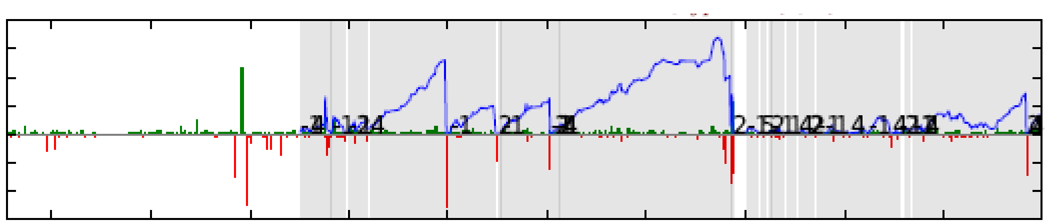

Sustained Balance – 1 Dump-and-pull – 11
Slow Drawdown – 1 Accumulator – 7
None – 11

Figure 20. Motifs Assigned for Account (c).

Account (c) has much activity, with sixty five segments over the one and a half years we have data for (Figure 20). (Similar to account (a), transaction data is missing for the first year). We can see some of the eight Accumulator patterns towering over the rest of the motifs. There are eleven Dump-and-Pulls, one each of Sustained Balance and Slow Drawdowns, and eleven segments which could not be mapped to a motif. A cursory inspection of this account would not suggest the balance reached zero or near-zero more than sixty times in the eighteen months; the motifs capture and classify this level of granularity comprehensively.

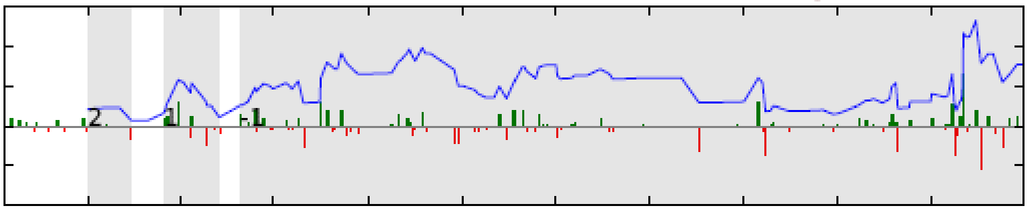

Figure 21. Motifs Assigned for Account (d).

Finally, account (d) has one Sustained Balance segment, one Dump-and-Pull segment, and a sprawling segment that could not be classified (Figure 21). Unlike previous examples where the Dump immediately followed the Pull, a certain amount of time elapses before we see the pull happening for this segment (labelled '2').

Motif Dominance

We conclude the process of pattern identification by introducing a concept called "motif dominance" that will be used heavily in subsequent analysis. We seek to devise an appropriate metric that captures overall motif behavior within a specific time frame. We recall that motifs are not particularly time-bound, in that a single motif could span months, and there could be multiple motifs within the same month. To capture this effect, we propose the concept of "motif dominance". This is motivated by the concept of "channel dominance" from the GAFIS project (BFA 2011), and is defined similarly:

An account is said to display **motif dominance** for a certain motif *if the number of segments mapping to that motif exceeds mapping to any other motif by at least 50%.*

The 50% threshold was chosen for “channel dominance” based on experiments on what number had the best discerning power. We simply port the concept to motif dominance and do not explore whether a different threshold could have worked better.

We use the same accounts we have been tracking throughout this thesis to illustrate this concept. The motif distributions are reproduced in Table 8 below.

ID	Sustained Balance	Dump-and-Pull	Fast Drawdown	Slow Drawdown	Accumulator	None
94e63e5576e8bf785d6b2b78702b800a	0	29	0	2	0	0
92982ecdee89aca7f4ed4bb20e156022	2	2	1	5	0	0
4305d4040a4d0cd136b073d070545006	1	11	0	1	7	11
7c5542fd93c457ade878f0edf720a005	1	1	0	0	0	1

Table 8. Motif Distribution for Illsutrative Accounts

For the first account, the 29 occurrences of Dump-and-Pulls are more than 50% as much again as the 2 occurrences of Slow Drawdowns; it can therefore be said to be “Fast Drawdown dominant”. For the second account, the 5 occurrences of Slow Drawdowns are more than 50% as much again as the 2 occurrences each of Sustained Balances and Dump-and-Pulls, and therefore can be labelled as “Slow Drawdown dominant.” Note that the definition requires the most frequent motif to pass the 50% test for the next numerous motifs, and not the total frequency of all other motifs. The third and fourth accounts do not have a dominant motif as there is a tie in terms of the highest motif frequencies.

Motif dominance can thus be considered to indicate behavioral preference over a given window of time.

Discussion of Motif Discovery Process

This chapter would not be complete without a discussion of the robustness of the motif discovery process we have detailed above. The entire process was borne out of a need to bucket time series account balance data in a meaningful way that would allow us to search for patterns. At the core of this process is the concept of a "segment," which recognizes the episodic nature of saving, and which we seek to bucket.

We are confident about the robustness of the process because it is the result of repeated iterations that ironed out various weaknesses, and it is fairly flexible in terms of how it deals with a dataset. The code that does the initial data processing has been iterated upon, tested and externally validated to the extent that we are confident in its accuracy. We do not prescribe fixed lengths of time over which savings patterns have to manifest themselves. Rather, we allow the zero-to-zero episodes to play out over whatever period of time it does so naturally. We also do not straitjacket the number of clusters or the number of data points needed in each segment, but rather arrive at them by seeking the numbers that give us the most reasonable configuration.

It is possible that a different dataset will require a different number of data points per segment, and a different number of clusters. This methodology is flexible enough to accommodate those differences easily.

While we cannot think of any specific vulnerability to this technique, we note some possibilities that may produce findings that are different from the five segment-based motifs:

- Considering a unit of analysis other than a "segment" as we have defined it may require different processing techniques. If one were to explore balance profiles in the frequency domain, as opposed to the time domain as we have done here, Fourier transforms may offer insights into other units of analysis. One could possibly construct other units of analysis by looking at "features" extracted from profiles, perhaps leveraging Bald's work (Bald 2008).

- Machine learning frameworks in the field of computer science have advanced tremendously. It is possible that feeding this data into those techniques would result in a typology of behavior expressed in parameters that are consistent with the learning framework.
- The accuracy of the Python code depends on its being able to handle all profile idiosyncrasies, and accounting for all necessary edge cases. Reasons this assumption may fail include: datasets containing edge cases not seen, conflating profile properties and the eternal scourge of coding – bugs!

Chapter 3: The Agent Difference

Identifying the Intervention

Now that we have our motifs, we explore whether it offers any additional explanatory power over simply looking at changes to balances and transactions as is. We noted how the induction of banking agents was accompanied by hopes that it would help clients interact more often with the bank, and save more. We also recognized that the reality would also be the opposite, facilitation of transactions leads to less savings because it makes it that much easier to withdraw funds. We will test those hypotheses in this section, in order to determine whether our five motifs can provide further details on what happens to client account usage patterns when they are introduced to accounts.

First, we use a simple comparison of balances, deposits and withdrawals before and after agent introduction within specific time windows to see if there are overarching trends. Then, we apply time series specific regression treatment in the form of Arellano Bond (AB) GMM Estimators to be rigorous and see if the findings hold. Finally, we see if we can make some kind of prediction based on what we see. We keep in mind that the entire purpose of this is to explore if motifs provide additional understanding on top of what we can already do.

Using simple means comparisons, we find that while account users deposit and withdraw more often after using agents for the first time, the change in average balances, and the average deposit and withdrawal amounts are not uniform. Accumulators and Sustained Balance accounts save more, Slow Drawdown accounts save less, the balances for Fast Drawdown accounts remain unchanged, and for Dump-and-Pull accounts, mostly unchanged. Fast Drawdowns and Accumulators to not change the amounts deposited after agent usage, while the rest decrease average deposit amounts. There is a general decrease in amounts withdrawn for all across motifs except Accumulators, which show no statistically significant change.

We will find that some of these relationships no longer persist, but confirm that in many cases, agent usage is associated with differentiated behavior based on motifs in subsequent sections through use of AB GMM Estimators.

Time Bound Comparisons of Means

Methodology

We would like to ascertain whether the frequencies and amounts of deposits and withdrawals, and balances, change after an agent has been used. A relatively straightforward technique to do this would be a t-test for each of these metrics before and after agent usage. This is essentially a t-test on a time dummy variable, and can be expressed as a regression specification:

$$y_i = \beta_0 + \beta_1 * D + u_{it} \qquad \text{(i)}$$

where:

$y \in$ {log(average balance), log(num deposits), log(num withdrawals), log(average deposit amount), log(average withdrawal amount)},

D = dummy variable, where 0 = before agent use, 1 = after agent use

y_i is the outcome variable – balances, deposits and withdrawals – for account *i*. *D* is the dumy variable that is set to 1 if the outcome variable is measured in the period after agent usage, and 0 if it is for the period before. If there is no difference in the outcome variables before and after agent usage, we will expect β_1 to be zero. β_0 represents the population mean of the outcome variable when there is no agent usage.

Since it is possible that balances can differ depending on what window of time we are looking at and we don't know exactly how long it takes for an "agent effect" to manifest itself, we test for 30-, 90-, 180- and 360-day windows. For example, for balances, this means that the average balance will be calculated for 30, 90, 180 and 360 days before or after the first agent transaction, depending on whether D is 0 or 1, respectively.

We are also interested to know if accounts that express different motif respond differently to the presence of an agent. We therefore run (i) above for one dominant motif at a time.

Recall that we introduced the concept of motif dominance at the end of Chapter 2, and defined it as the motif that expresses itself at least 50% more often as the next frequent motif within a given period of time. This allows us to explore if motifs might have explanatory power on top of agent usage, if we do indeed see a difference in balances. Note that simply using dominant motif as a parameter in (ii) below is insufficient to explore this question, as that simply tells us whether β_0 needs be adjusted for particular motifs and nothing about how β_1 could be affected.

$$y_{\mathrm{i}} = \beta_0 + \beta_1 * D + \beta_2 * dominant\_motif_{\mathrm{i}} + u_{\mathrm{it}} \tag{ii}$$

The approach of running (i) on a per-dominant motif level has its own limitations, as discussed in the section *The Need to Delve Further*, but it is a decent start to getting a sense of possible relationships between motifs and agent usage.

We will also be using the logarithm of all the outcome variables, as they all resemble an exponential decay function, but assume distributions that are relatively normal once the log-transform is applied. This will be demonstrated in the relevant sections.

Comparing Balances Before and After Agent Use

This section explores whether balance values increase or decrease after an account holder uses an agent. First, we compare the average balances before and after the agent usage for the entire portfolio, and then by behavioral motifs.

The exponential decay shape of raw balance values is evident in . The transformation through natural logarithms of these balances does provide a normal distribution, as seen in . We leave out the largest 1% of values from the untransformed values in to allow us to see the shape of the remaining 99%.

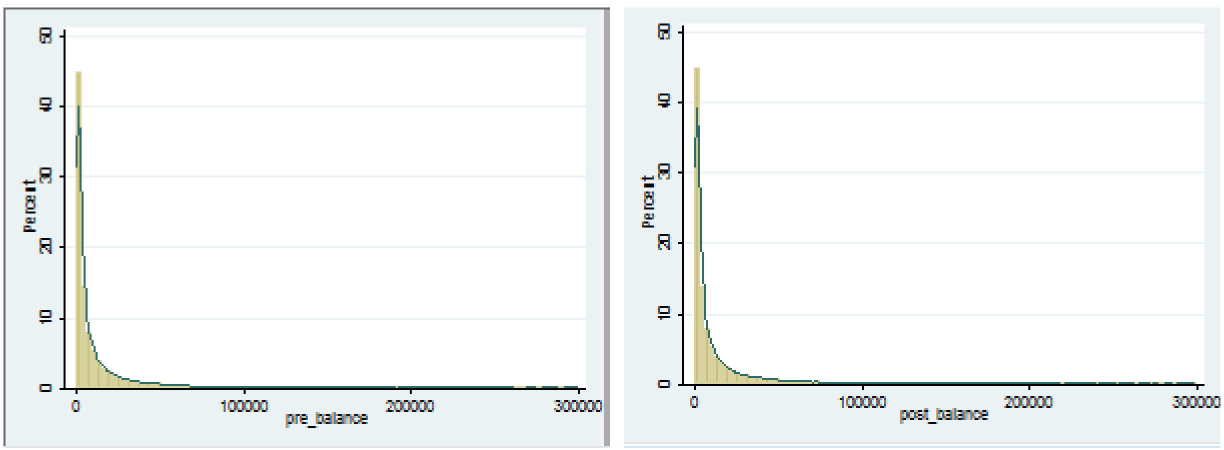

Figure 22. Balances Before (Left) and After (Right) Agent Usage

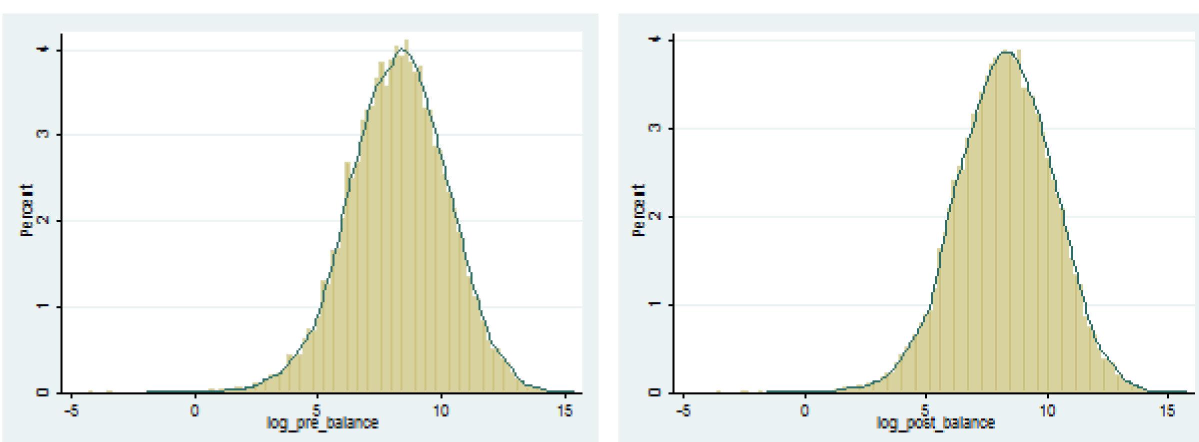

Figure 23. Natural Logarithm of Balances Before (Left) and After (Right) Agent Usage

The results for the regression with the time dummy variable only (i) is presented in Table 9 below. None of the p-values β_1 are near 0.10, which would imply significance at a 10% level, let alone near 0.05 or 0.01 (i.e. significance levels of 5% or 1% respectively). This tells us that, at first blush, there is no correlation between agent usage and changing balances within a one-year time frame.

Days	# Accounts	β_0	P-value for β_0	β_1	P-value for β_1
30	25,836	8.1618	0.0000	(0.0203)	0.5234
90	22,862	8.2627	0.0000	(0.0020)	0.9456
180	20,034	8.3489	0.0000	(0.0027)	0.9271
360	15,396	8.5146	0.0000	(0.0222)	0.4707

Table 9. Time-dummy regression results for average balance

Now, we conduct the same exploration using motif dominance. If they have no discerning power, we would expect accounts to follow the overall trend as outlined in Table 9 – that averages balances do not change after agent usage. The results for this are presented in Table 10 below, with coefficients accompanied by their p-values. The sample size reduces markedly as we take larger and larger windows on each side of the first agent usage.

We only run regressions if data is available for the full duration of the 30- to 360-day window we are interested in to prevent short term balance anomalies from skewing results through an outlier effect. We also only run regressions if the account displays the same dominant motif behavior before and after agent usage. These are limiting conditions that will be addressed in subsequent sections.

Dominant Motif	Days	# Accounts	Population Mean, Average Balance (β_0)	p-value for β_0	Change Associated with Agent Usage (β_1)	p-value for β_1
Sustained Balance	30	4,268	8.7879	0.0000	0.1320	0.0379
	90	3,081	8.7781	0.0000	0.1684	0.0236
	180	2,304	8.8733	0.0000	0.1502	0.0756
	360	1,467	9.0831	0.0000	0.0500	0.5914
Dump and Pull	30	4,414	7.4398	0.0000	(0.1251)	0.1250
	90	5,100	7.7669	0.0000	(0.0636)	0.3230
	180	5,269	7.9942	0.0000	(0.1130)	0.0491
	360	4,584	8.2120	0.0000	(0.1813)	0.0021
Fast Drawdown	30	911	6.3823	0.0000	(0.1112)	0.7135
	90	751	6.6423	0.0000	(0.2172)	0.4448
	180	566	6.6768	0.0000	0.0967	0.7508
	360	332	7.2015	0.0000	0.1742	0.4456
Accumulator	30	2,657	7.6597	0.0000	0.2958	0.0304
	90	1,733	7.9394	0.0000	0.5531	0.0001
	180	1,251	7.8854	0.0000	0.8098	0.0000
	360	797	8.2049	0.0000	0.7796	0.0000
Slow Drawdown	30	3,663	8.6386	0.0000	(0.1757)	0.0053
	90	2,926	8.5182	0.0000	(0.2127)	0.0009
	180	2,375	8.5757	0.0000	(0.2341)	0.0005
	360	1,654	8.6823	0.0000	(0.2701)	0.0003
No Dominant Motif	30	5,828	7.9138	0.0000	(0.1585)	0.0014
	90	6,209	8.1600	0.0000	(0.0713)	0.1080
	180	5,925	8.2722	0.0000	(0.0422)	0.3550
	360	5,016	8.3671	0.0000	0.1431	0.0033

Table 10. Time-dummy Regression Results for Average Balance and Dominant Motifs

It seems that balances change with agent behavior in different ways, depending on what dominant motif they are. Accumulator and Sustained Balance dominant accounts show an increase in balances, Slow Drawdown dominant accounts show a decrease in balances, and Fast Drawdown accounts show no change in balances at a statistically significant level, after agent usage, for all four time windows. Dump-and-Pull dominant accounts show a slight decrease in balances for the longest time windows.

Some of the balance changes we see make intuitive sense. Accumulator accounts represent accretive savings, and the increase in balances could be a result of the fact that agents make it easier to save in those smaller amounts compared to balances, as predicted by literature. The Slow Drawdown pattern represent a case where an initial funds injection is drawn down at different rates over time, and agents may be facilitating more withdrawals, thus leading credence to the concern we observed in literature about agents not facilitating savings behavior because they make withdrawing easier. Dump-and-pull accounts understandably show no change in balances as they withdraw everything, or almost everything, thereby leaving very little scope for significant changes in balance.

Thus, **banking agents seem to reinforce whatever accounts are prone to do already, and not cause fundamental changes in behavior**.

Differentiated behavior along motifs suggests that there is value in including motifs in the analysis of outcome variables, as opposed to simply using raw data. This simple, sub-sample regression underlines the necessity of a more rigorous analysis, as it is possible that whatever behavioral changes resulted in changed balances could also result in a change in motif expression. Also, given that we see such changes in balances around the very first agent usage, it is possible that differences are reinforced with continued agent usage. We need to be able to account for continual motif expressions and agent usage.

Before we look at addressing these issues, let us apply the same treatment for deposit and withdrawal counts and amounts.

Comparing Transaction Counts Before and After Agent Usage

We replicate the same treatment with transactions initiated by the client, as we did with balances in the previous section. We consider both the number and amount of customer-initiated (CI) transactions, and consider deposits and withdrawals separately. We ignore fees and other business-initiated (BI) transactions, as they usually accompany CI transactions and counting them would erroneously amplify the number of transactions clients are conducting.

We take the number of deposits and withdrawals on a per thirty day basis, for any given time window. This is done to keep the metrics comparable across time windows, as we have data for some accounts for a month only, while we have data for more than a year for others, and taking the raw transaction count would mean the latter group would register figures that are an order of magnitude more than the former. This interpolation implies that if we have *n* days of transaction data, the 30-day transaction count is given by:

$$\text{Transaction count, 30-days} = (\text{Transaction count, raw}) \times 30.0 / n \qquad \text{(iii)}$$

We contend that we must use the natural logarithm for the number of deposits and withdrawals that take place, as they too display strong positive skew and the transformation provides us a more normal distribution, as can be seen in Figure 24.

Raw Values

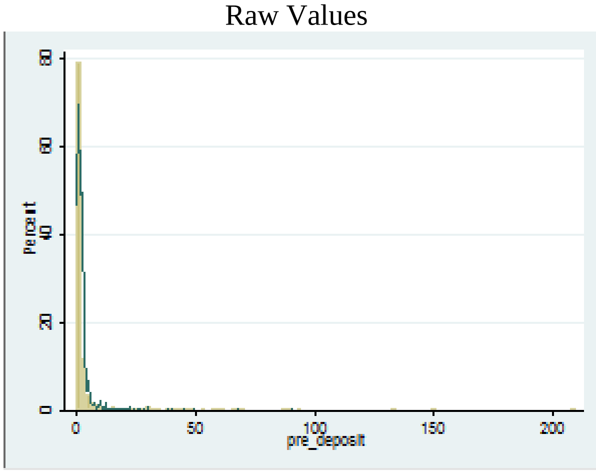

Deposits, Before

Natural Logarithm of Raw Values

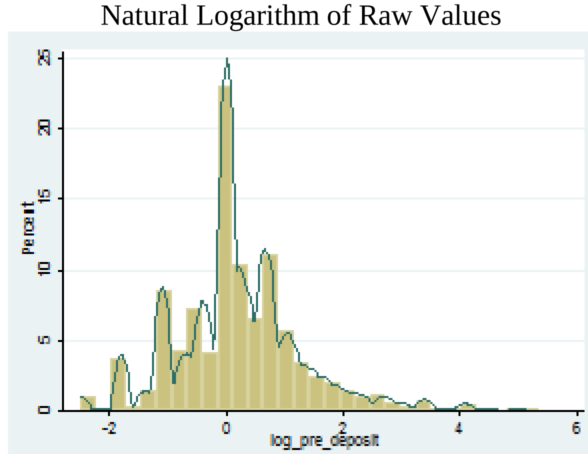

log(# Deposits), Before

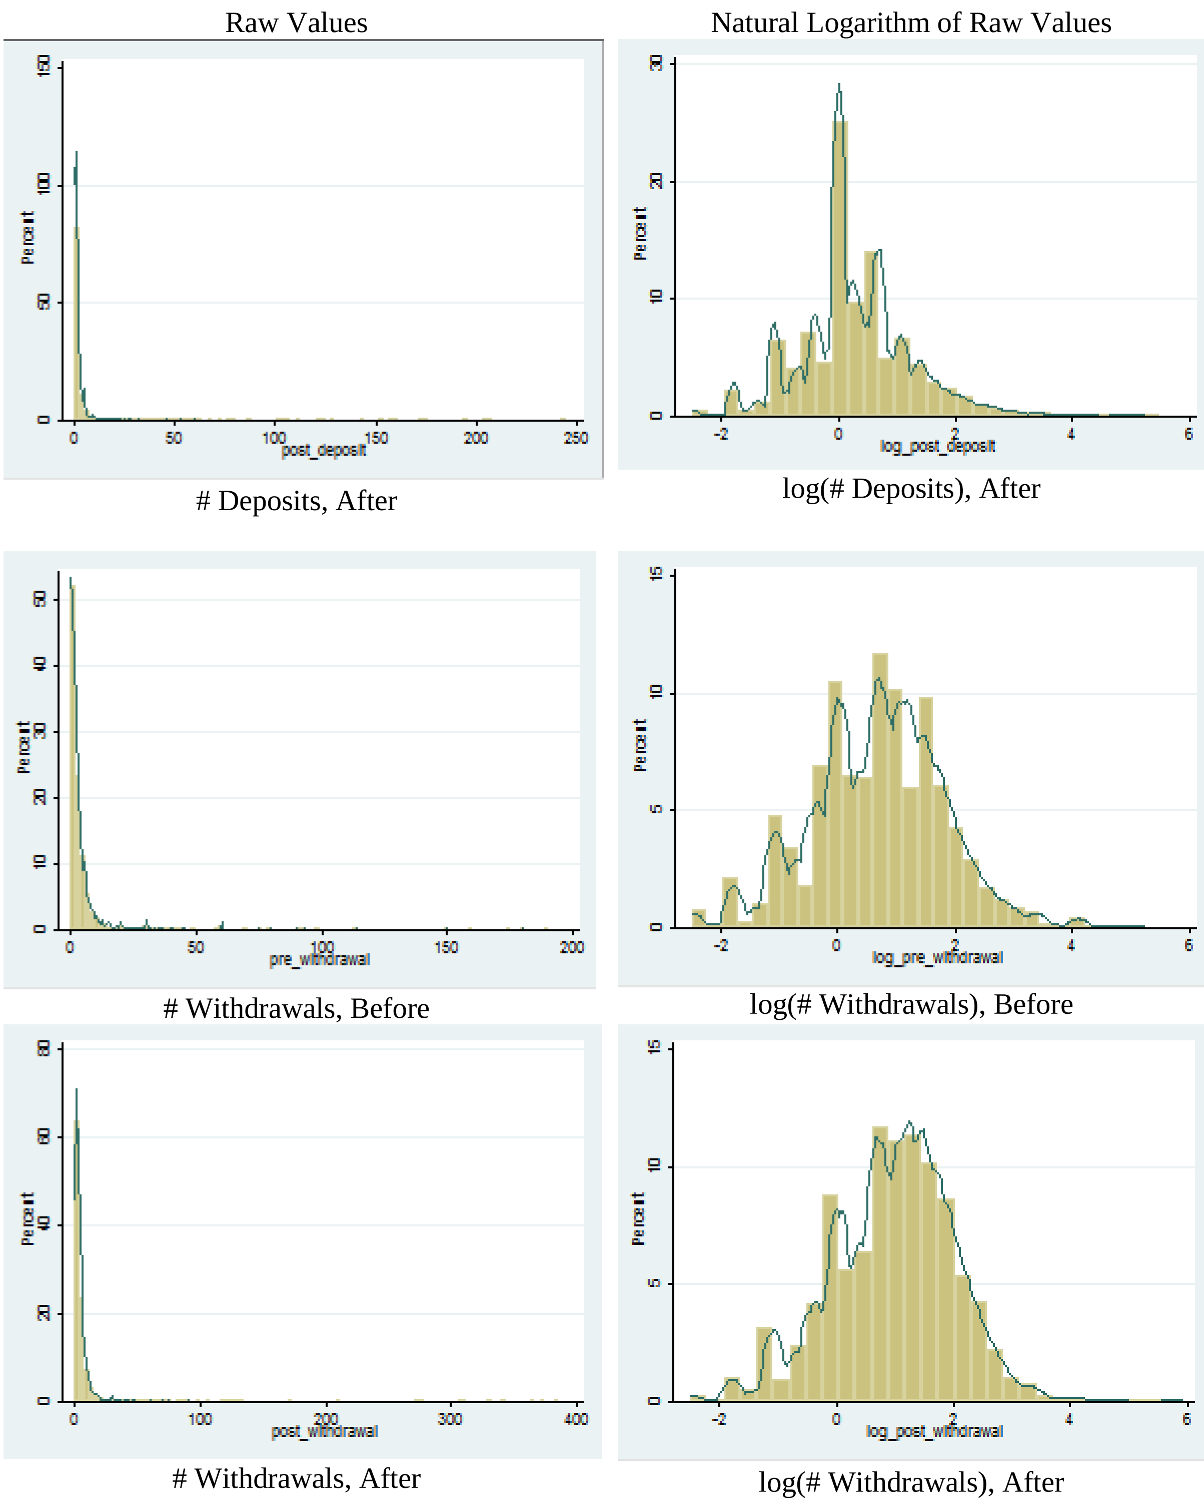

Figure 24. Log Transformation of Number of Deposits and Withdrawals

As before, we are looking at 30-, 90-, 180- and 360-day windows. The time-dummy regression results for the transformed values of deposits and withdrawals are presented in Table 11 and Table 12 below.

Days	# Accounts	β_0	P-value for β_0	β_1	P-value for β_1
30	21,608	0.4567	0.0000	0.1907	0.0000
90	21,308	(0.0304)	0.0001	0.2469	0.0000
180	19,382	(0.2408)	0.0000	0.3127	0.0000
360	15,207	(0.3250)	0.0000	0.3535	0.0000

Table 11. Time-dummy Regression Results for log(Number of Deposits)

Days	# Accounts	β_0	P-value for β_0	β_1	P-value for β_1
30	22,288	1.0772	0.0000	0.2073	0.0000
90	20,985	0.6171	0.0000	0.2749	0.0000
180	19,001	0.4243	0.0000	0.3238	0.0000
360	15,022	0.3175	0.0000	0.3736	0.0000

Table 12. Time-dummy Regression Results for log(Number of Withdrawals)

The p-values for deposits and withdrawals are 0.0000, which implies that these deposit and withdrawal counts are statistically significantly different before and after first agent usage. Both seem to *increase* in occurrence after the first agent usage, as illustrated by the p-value of β_1. There is a rather sharp drop in β_0 as the time-window increases. Given the diminishing sample size, it's not clear if this is because of some inherent changes of accounts over time, or if it is an artefact of the sample itself. The diminishing sample size is a result of the constraint that we only consider the normalized transaction counts if at least as many days of transaction exists as the window being considered.

Do note that the very strong association for deposits and withdrawals cannot be taken to be causal because other possible explanatory variables are not considered (and cannot be considered because we don't have additional data for that purpose), but given the large sample size, we can say that when clients use an agent, we can expect transaction counts to increase.

Now, we introduce motif dominance and see if the pattern of increased deposit counts holds across motifs. The results for this are presented in Table 13 below.

Dominant Motif	Days	# Accounts	Population Mean, Average Balance (β_0)	p-value for β_0	Change Associated with Agent Usage (β_1)	p-value for β_1
Sustained Balance	30	3,321	0.4674	0.0000	0.2212	0.0000
	90	2,767	(0.1374)	0.0000	0.3568	0.0000
	180	2,193	(0.4174)	0.0000	0.4333	0.0000
	360	1,446	(0.5421)	0.0000	0.4743	0.0000
Dump and Pull	30	3,638	0.4166	0.0000	0.2099	0.0000
	90	4,735	0.0842	0.0000	0.1588	0.0000
	180	5,115	(0.0446)	0.0180	0.1488	0.0000
	360	4,527	(0.0917)	0.0000	0.1192	0.0000
Fast Drawdown	30	708	0.2013	0.0000	0.0785	0.0169
	90	690	(0.1687)	0.0000	0.1517	0.0014
	180	546	(0.3517)	0.0000	0.1829	0.0035
	360	330	(0.4872)	0.0000	0.2994	0.0012
Accumulator	30	2,317	0.3928	0.0000	0.3767	0.0000
	90	1,697	(0.1832)	0.0000	0.5062	0.0000
	180	1,243	(0.4415)	0.0000	0.6642	0.0000
	360	798	(0.5149)	0.0000	0.7315	0.0000
Slow Drawdown	30	2,580	0.3019	0.0000	0.1358	0.0000
	90	2,517	(0.1967)	0.0000	0.1865	0.0000
	180	2,209	(0.4225)	0.0000	0.2552	0.0000
	360	1,607	(0.5131)	0.0000	0.2206	0.0000
No Dominant Motif	30	5,518	0.6289	0.0000	0.0351	0.0655
	90	6,018	0.0878	0.0000	0.1468	0.0000
	180	5,811	(0.1771)	0.0000	0.2621	0.0000
	360	4,976	(0.3337)	0.0000	0.4363	0.0000

Table 13. Time-dummy Regression Results for log(Deposit Counts), by Dominant Motifs

The number of deposits has increased across the board, across every single time window, as seen by the statistically significant values of β_1. Accumulators and Sustained Balances have the highest increase in deposits, while Fast Drawdowns have the smallest increase.

Now, we introduce motif dominance and see if the pattern of increased deposit counts holds across motifs. The results for this are presented in Table 14 below.

Dominant Motif	Days	# Accounts	Population Mean, Average Balance (β_0)	p-value for β_0	Change Associated with Agent Usage (β_1)	p-value for β_1
Sustained Balance	30	3,621	1.1079	0.0000	0.2126	0.0000
	90	2,783	0.5864	0.0000	0.3225	0.0000
	180	2,152	0.3611	0.0000	0.3568	0.0000
	360	1,413	0.2179	0.0000	0.3603	0.0000
Dump and Pull	30	3,779	0.8359	0.0000	0.2397	0.0000
	90	4,749	0.6160	0.0000	0.2311	0.0000
	180	5,078	0.4889	0.0000	0.2772	0.0000
	360	4,514	0.4347	0.0000	0.2146	0.0000
Fast Drawdown	30	849	0.9345	0.0000	0.1967	0.0000
	90	715	0.6967	0.0000	0.1450	0.0226
	180	545	0.5407	0.0000	0.1757	0.0238
	360	326	0.3975	0.0000	0.3587	0.0008
Accumulator	30	1,775	0.5732	0.0000	0.1665	0.0000
	90	1,242	0.0969	0.0135	0.1968	0.0006
	180	962	(0.1545)	0.0031	0.2500	0.0010
	360	682	(0.3425)	0.0000	0.3703	0.0003
Slow Drawdown	30	3,440	1.2679	0.0000	0.1947	0.0000
	90	2,821	0.8722	0.0000	0.2692	0.0000
	180	2,313	0.7541	0.0000	0.3150	0.0000
	360	1,632	0.6786	0.0000	0.3528	0.0000
No Dominant Motif	30	5,753	1.3769	0.0000	0.1703	0.0000
	90	6,133	0.7281	0.0000	0.3137	0.0000
	180	5,859	0.4793	0.0000	0.3656	0.0000
	360	4,986	0.3371	0.0000	0.5215	0.0000

Table 14. Time-dummy Regression Results for log(Withdrawal Counts), by Dominant Motifs

Thus, it seems that withdrawal counts increase at statistically significant levels across all motifs. The highest increase is with Sustained Balances in general. The increase in withdrawal frequency is higher across larger time windows.

Based on these time-dummy regressions that take dominant motifs into account, we can conclude that there is evidence that an increase in both deposit and withdrawal frequencies is associated with agent usage. Even though there isn't differentiated behavior along motifs when it comes to transaction frequencies as far as the direction of change is concerned, the magnitude of change is different across motifs, suggesting that there is continued value in including motifs in the analysis of outcome variables, as opposed to simply using raw data.

Comparing Transaction Amounts Before and After Agent Usage

We replicate the same treatment with average transaction amounts initiated by the client, as we did with transaction counts in the previous section. The average deposit and withdrawal amounts are calculated by simply summing the total amount of transactions of that type and dividing by the number of said transactions, for a given time window.

We use the natural logarithm for the average amounts of deposits and withdrawals that take place too, as they too display strong positive skew and the transformation provides us a more normal distribution over which means-comparison t-tests can be conducted Figure 25. We leave out the largest 1% of values from the untransformed values to allow us to see the shape of the remaining 99%. Note that withdrawn amounts are negative, hence we take the natural logarithm of the absolute value as negative values have no logarithms.

Raw Values

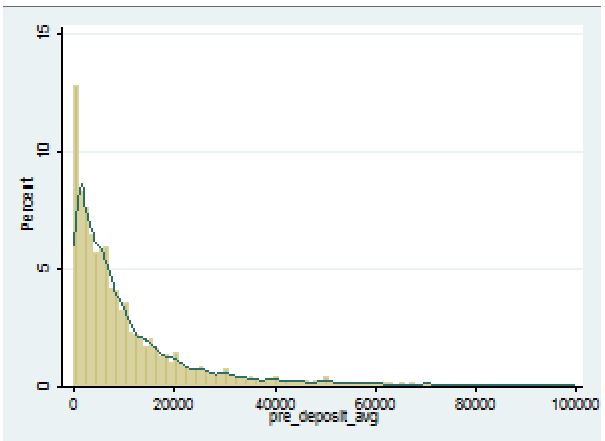

Average Deposit Amount, Before

Natural Logarithm of Raw Values

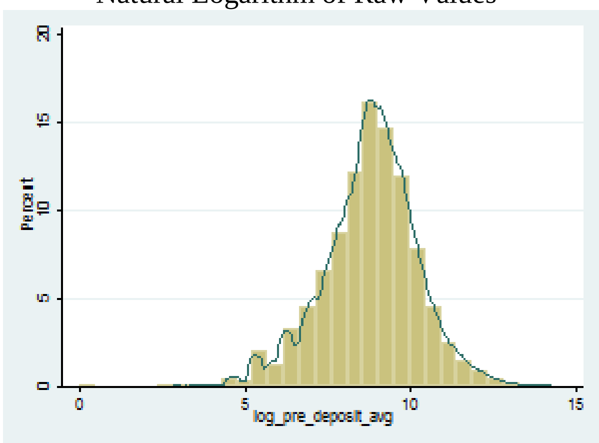

log(Average Deposit Amount), Before

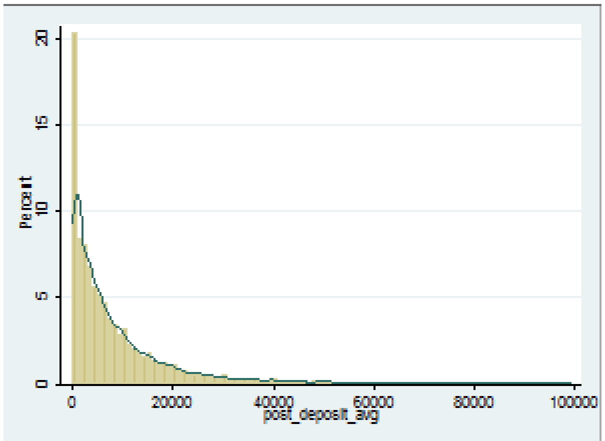

Average Deposit Amount, After

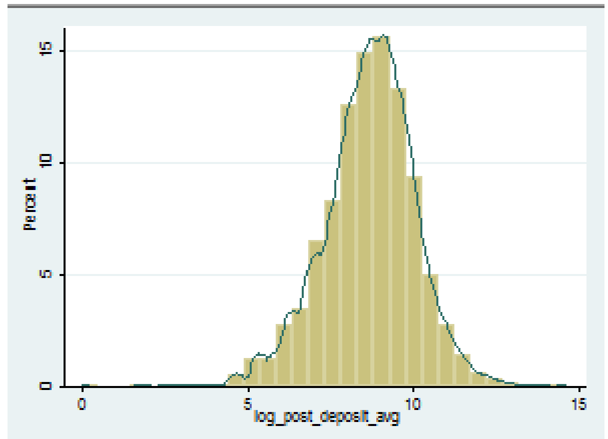

log(Average Deposit Amount), After

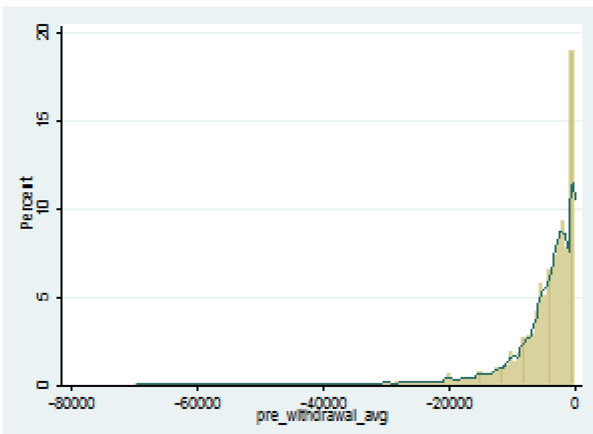

Average Withdrawal Amount, Before

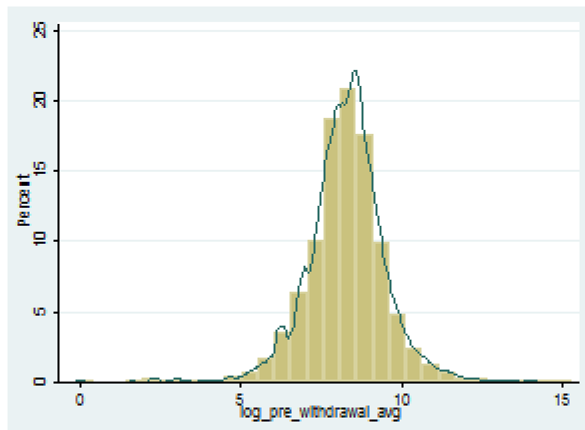

log(Average Withdrawal Amount), Before

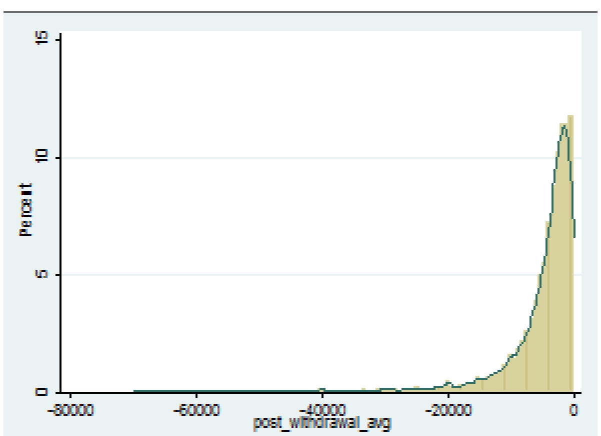

Average Withdrawal Amount, After

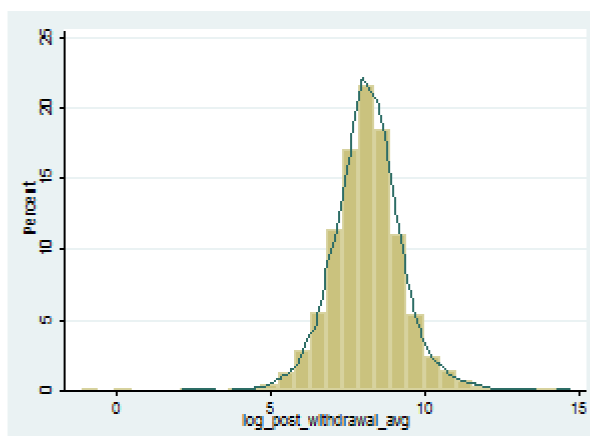

log(Average Withdrawal Amount), After

Figure 25. Log Transformation of Deposit and Withdrawal Average Amounts

As before, we are looking at 30-, 90-, 180- and 360-day windows. The time-dummy regression results for the transformed values of deposits and withdrawals are presented in Table 15 and Table 16 below, along with the untransformed values.

Days	# Accounts	β_0	P-value for β_0	β_1	P-value for β_1
30	21,608	8.6920	0.0000	(0.1606)	0.0000
90	21,308	8.8669	0.0000	(0.1546)	0.0000
180	19,382	8.9834	0.0000	(0.1567)	0.0000
360	15,207	9.1403	0.0000	(0.1732)	0.0000

Table 15. Time-dummy Regression Results for log(Average Deposit Amount)

Days	# Accounts	β_0	P-value for β_0	β_1	P-value for β_1
30	22,288	8.1421	0.0000	(0.1191)	0.0000
90	20,985	8.3259	0.0000	(0.1743)	0.0000
180	19,001	8.4366	0.0000	(0.1893)	0.0000
360	15,022	8.5429	0.0000	(0.1737)	0.0000

Table 16. Time-dummy Regression Results for log(Average Withdrawal Amount)

All values for β_1 are negative, and all corresponding p-values are 0.0000, which implies that both the average deposit and withdrawal amounts are statistically significantly *lower* after first agent usage, compared to before.

Now, we introduce motif dominance and see if the pattern of decreased deposit amounts holds across motifs. The results for this are presented in Table 17.

Dominant Motif	Days	# Accounts	Population Mean, Average Balance (β_0)	p-value for β_0	Change Associated with Agent Usage (β_1)	p-value for β_1
Sustained Balance	30	3,321	8.7020	0.0000	(0.1918)	0.0001
	90	2,767	8.8133	0.0000	(0.1276)	0.0149
	180	2,193	8.9078	0.0000	(0.1707)	0.0024
	360	1,446	9.0378	0.0000	(0.2562)	0.0001
Dump and Pull	30	3,638	8.3907	0.0000	0.0005	0.9923
	90	4,735	8.7748	0.0000	(0.0464)	0.2430
	180	5,115	8.9906	0.0000	(0.1369)	0.0002
	360	4,527	9.1914	0.0000	(0.2357)	0.0000
Fast Drawdown	30	708	8.8068	0.0000	(0.1255)	0.2218
	90	690	9.0324	0.0000	(0.1250)	0.1447
	180	546	9.0929	0.0000	(0.1151)	0.1960
	360	330	9.2023	0.0000	(0.1344)	0.2131
Accumulator	30	2,317	8.1634	0.0000	(0.2012)	0.0012
	90	1,697	8.2214	0.0000	(0.0722)	0.3224
	180	1,243	8.3278	0.0000	(0.0021)	0.9798
	360	798	8.4923	0.0000	(0.0086)	0.9278
Slow Drawdown	30	2,580	9.2599	0.0000	(0.4364)	0.0000
	90	2,517	9.4328	0.0000	(0.3319)	0.0000
	180	2,209	9.4909	0.0000	(0.2364)	0.0000
	360	1,607	9.5532	0.0000	(0.1150)	0.0276
No Dominant Motif	30	5,518	9.0895	0.0000	(0.2567)	0.0000
	90	6,018	9.0883	0.0000	(0.2534)	0.0000
	180	5,811	9.1210	0.0000	(0.2330)	0.0000
	360	4,976	9.2076	0.0000	(0.1379)	0.0001

Table 17. Average Deposit Amounts after Agent Usage, by Motif Dominance.

The trend of lower deposit amounts is evident for accounts that are Slow Drawdown, Sustained Balance, or have no dominant motif. Accounts that are Fast Drawdown have no statistical change in average deposit amounts. The average deposit value of Accumulators decreases in the 30-day window, but is unchanged for the other three. Dump-and-Pull accounts are the opposite, with there being no change in the shorter time windows, but there being a decrease in the 180- and 360-day windows. The general tendency then is for Fast Drawdowns and Accumulators to not change the amounts deposited after agent usage, while the rest decrease average deposit amounts.

Finally, we conduct the same analysis for average withdrawal amounts. The results for this are presented in Table 18.

Dominant Motif	Days	# Accounts	Population Mean, Average Balance (β_0)	p-value for β_0	Change Associated with Agent Usage (β_1)	p-value for β_1
Sustained Balance	30	3,621	8.0622	0.0000	(0.1279)	0.0017
	90	2,783	8.2059	0.0000	(0.1256)	0.0041
	180	2,152	8.2931	0.0000	(0.1646)	0.0005
	360	1,413	8.3888	0.0000	(0.1481)	0.0078
Dump and Pull	30	3,779	8.0836	0.0000	(0.0446)	0.2777
	90	4,749	8.4161	0.0000	(0.1731)	0.0000
	180	5,078	8.5828	0.0000	(0.2346)	0.0000
	360	4,514	8.7134	0.0000	(0.2728)	0.0000
Fast Drawdown	30	849	8.1418	0.0000	(0.2880)	0.0002
	90	715	8.3029	0.0000	(0.3495)	0.0000
	180	545	8.3768	0.0000	(0.3584)	0.0000
	360	326	8.4366	0.0000	(0.2574)	0.0025
Accumulator	30	1,775	8.1781	0.0000	(0.0270)	0.6515
	90	1,242	8.2204	0.0000	0.0137	0.8606
	180	962	8.3898	0.0000	(0.0070)	0.9329
	360	682	8.4836	0.0000	0.0036	0.9698
Slow Drawdown	30	3,440	8.1548	0.0000	(0.2540)	0.0000
	90	2,821	8.2770	0.0000	(0.2895)	0.0000
	180	2,313	8.3671	0.0000	(0.3072)	0.0000
	360	1,632	8.3811	0.0000	(0.2363)	0.0000
No Dominant Motif	30	5,753	8.2641	0.0000	(0.1523)	0.0000
	90	6,133	8.3818	0.0000	(0.2034)	0.0000
	180	5,859	8.4192	0.0000	(0.1715)	0.0000
	360	4,986	8.5151	0.0000	(0.1208)	0.0000

Table 18. Average Withdrawal Amounts after Agent Usage, by Motif Dominance.

The general decrease in withdrawal amounts is evident for all motifs except Accumulators, which show no statistically significant change. Dump-and-pull accounts show no change in the 30-day window, but decrease across all other windows. Thus, we can say that the overall pattern of decreased average withdrawal amounts holds across all motifs, except Accumulators.

Incidentally, the differentiated behavior along motifs when it comes to transaction amounts across motifs suggests that there is continued value in including motifs in the analysis of these outcome variables, as opposed to simply using raw data.

Combined Impact on Balances and Transactions

We combine the information presented in the previous sections to come up with an overview of how account usage behavior changes for the four windows, given a certain pre-dominant motif. Only difference in pre- and post- agent values which are significant at a 0.10 level or lower are noted in Table 19. We are looking for patterns that set the motifs apart.

Pre-Dominant Motif	Days	Δ log (Average Balance)	Δ log (Deposit Count)	Δ log (Withdrawal Count)	Δ log (Deposit Amount)	Δ log (Withdrawal Amount)
Sustained Balance	30	0.13	0.22	0.21	(0.19)	(0.13)
	90	0.17	0.36	0.32	(0.13)	(0.13)
	180	0.15	0.43	0.36	(0.17)	(0.16)
	360		0.47	0.36	(0.26)	(0.15)
Dump and Pull	30		0.21	0.24		
	90		0.16	0.23		(0.17)
	180	(0.11)	0.15	0.28	(0.14)	(0.23)
	360	(0.18)	0.12	0.21	(0.24)	(0.27)
Fast Drawdown	30		0.08	0.20		(0.29)
	90		0.15	0.15		(0.35)
	180		0.18	0.18		(0.36)
	360		0.30	0.36		(0.26)
Accumulator	30	0.30	0.38	0.17	(0.20)	
	90	0.55	0.51	0.20		
	180	0.81	0.66	0.25		
	360	0.78	0.73	0.37		
Slow Drawdown	30	(0.18)	0.14	0.19	(0.44)	(0.25)
	90	(0.21)	0.19	0.27	(0.33)	(0.29)
	180	(0.23)	0.26	0.31	(0.24)	(0.31)
	360	(0.27)	0.22	0.35	(0.11)	(0.24)
No Dominant Motif	30	(0.16)	0.04	0.17	(0.26)	(0.15)
	90		0.15	0.31	(0.25)	(0.20)
	180		0.26	0.37	(0.23)	(0.17)
	360	0.14	0.44	0.52	(0.14)	(0.12)

Table 19. Change in Balance and Transaction Metrics after Agent Usage

We make the following observations in order of most important findings as related to differences in average balances, the number and average size of deposits, and the number and average size of withdrawals:

- Both the number of deposits and withdrawals undertaken increase after the first agent use for every dominant motif, as well as accounts which have no dominant motif.

- **Accumulators** show higher balances after agent usage across all time windows, but this is not accompanied by higher savings amounts or smaller withdrawal amounts. In so far as the change in deposit frequency is twice that of the change in withdrawal frequency, we can hypothesize that the higher balances are a result of similar amounts being deposited more frequently. These could be cashflow positive individuals who are taking advantage of the proximity of agents to deposit cash at their convenience. Agents seem to empower these accounts which are deliberate, accretive savers, to do more of the same.
- **Fast Drawdown** accounts have no change in balances, but do show a decrease in the average withdrawal amounts. There is no change in the deposit amounts. This seems to be a case where the initial deposit hasn't changed perhaps because of the nature of the source, but access to agents gives the client agency in not having to draw as much down as they used to before after the initial large drawdown is made, allowing them to accumulate more funds on average over time. Whatever the reason may be, agents may thus be allowing these accounts to save more of their residual funds.
- **Slow Drawdown** accounts show the opposite movement for average balance. Both deposit and withdrawal amounts are smaller, but the decrease in deposit amounts is more than the decrease in withdrawal amounts, causing average balances to fall for all time windows. It is unclear as to what behavioral pattern this might fit with for these motifs we identified with as being used to fund regular expenses.
- **Sustained Balance** accounts show an increase in balances in three out of four windows, and both deposit and withdrawal amounts decrease after agent usage. These accounts holders thus transact more in smaller amounts, but the net result in terms of balances increases somewhat.
- **Dump-and-Pull** accounts do not show any change in balance in the shorter time windows, but do show a decrease in the 180- and 360day windows. This is accompanied by a decrease in deposit amounts in these windows, which suggests that in the long run, these accounts may

end up depositing less, and therefore having a lower balance amount. The lack of changes in balances and amounts is consistent with what we intuitively understand the pattern to be, as everything or almost everything is withdrawn in the first incidence after a deposit, so the chance to alter average balances is slim.

- Accounts which could not be assigned a dominant motif show changes in balances in both direction, as well as no change, depending on the time window. Deposit and Withdrawal amounts consistently go down.

Overall, a case can be made that **agent usage is associated with reinforcement of existing behavior** in some cases. Accumulators save even more, driven by a greater frequency of deposits. Sustained Balance accounts also increase their average balance levels, exercising greater granularity in how often and in what quantities funds are intermediated. In so far as Fast Drawdowns represent attempts at residual savings after an initial large withdrawal, a reduction in withdrawal amounts suggests a greater amount of such residuals. And since Dump-and-Pulls have no residual balance to change, we are not surprised to see no change in balances or transaction amounts in the near term after agent usage. Slow Drawdowns are the exception to this trend of "business as usual, simply with more vigor," as we cannot explain why deposit amounts are reduced by a significant amount, resulting in lower average balances.

With these initial insights in mind, we move on to a more rigorous treatment of balance and transactional changes in the next section, to navigate the morass of challenges that time series data are wont to offer.

The Need to Delve Further

A simple means comparison test as presented earlier does not capture the full dynamics of what might be happening, as there are multiple threats to identification that remain unaccounted for. The model that was tested is, initially across the entire portfolio, and then for each dominant motif at a time:

$$y_i = \beta_0 + \beta_1 * D + u_{it}$$

where:

$y \in \{log(average\ balance), log(num\ deposits), log(num\ withdrawals),$
$log(average\ deposit\ amount), log(average\ withdrawal\ amount)\}$

There are a couple of shortcomings of this approach of running a time-dummy based regression for the same dominant motif, comparing a fixed time period before and after the first agent usage. We know that motifs can express themselves over arbitrary amounts of time. As a result some 90-day windows might have a single motif, while others might have a handful. It would be preferable to track outcome variables within the same segment change over time. We also want to be able to account for motifs in the regression itself without breaking the population into sub-samples. We can achieve both by using the following specification:

$$y_{i,t} = \beta_0 + \beta_1 * motif_{i,t} + \beta_2 * used\_agent_{i,t} + u_{it}$$

$y_{i,t}$ is the outcome variables – balances, deposits and withdrawals – at time t, for account i. $used\_agent_{i,t}$ is our exogenous independent variable, and $motif_{i,t}$ is our potentially endogenous independent variable.

Note that we have included *motif* as an explicit explanatory variable in the equation as we are no longer restricting comparison runs by dominant motifs. Because there are five dominant motifs and a sixth "no dominant motif" category, *motif* is actually a vector of six elements, each of which is a dummy variable for the corresponding motif type:

$$\overrightarrow{motif} = \begin{bmatrix} Motif\ 1\ dummy\ \in \{0,1\} \\ Motif\ 2\ dummy\ \in \{0,1\} \\ Motif\ 3\ dummy\ \in \{0,1\} \\ Motif\ 4\ dummy\ \in \{0,1\} \\ Motif\ 5\ dummy\ \in \{0,1\} \\ No\ motif\ dummy\ \in \{0,1\} \end{bmatrix}$$

Another shortcoming of the simple specification is that motifs and agent usage are unlikely to be the only factors explaining balance levels or the number or amount of transactions. The gender of the account user, whether they live in rural or urban areas, what they do for a living, distance to nearest agent, etc. could all influence how much they earn, how much they decide to save, and how often they can interact with agents. These **omitted variables** could be time-invariant, such as gender and location, in which case they are called fixed effects, or time-variant, such as distance to agent, which evolve over time. Our dataset has no information on such factors, potentially leaving us with omitted variable bias. We will have more to say about time-variant omitted variables in due course; for now, we expand our specification to include the fixed effects, a_i. Note that the time subscript, *t*, is absent from a_i.

$$y_{i,t} = \beta_0 + \beta_1 * motif_{i,t} + \beta_2 * used\_agent_{i,t} + \beta_3 * a_i + u_{it}$$

Another factor that could help effect the outcome variable – balances, deposits and withdrawals – can depend on its own past realizations. How much and how often individuals save and spend is determined by the circumstances of their financial involvements. Incomes can be regular, from a salaried job, or random, from a one-time sale of assets. Similarly, expenses can be regular, such as food expenses, or random, such as unexpected illness. To the extent that these accounts are topped up with income from regular sources and used to pay regular expenses, there will be correlation between the amount and frequency of inflows and outflows of adjacent periods. Given that the balance at time *t* is obtained from the balance at time *t-1* after netting deposits and withdrawals, as well as the fact a certain balance is maintained deliberately or through momentum, we can also expect balances to be correlated across time periods. We thus include the **lagged outcome variable**, $y_{i,t-1}$, in our specification. Incidentally, the presence of a lagged outcome variable characterizes the panel data as "dynamic".:

$$y_{i,t} = \beta_0 + \beta_1 * motif_{i,t} + \beta_2 * used\_agent_{i,t} + \beta_3 * y_{i,t-1} + \beta_4 * a_i + u_{it}$$

We propose to add the lagged motif, $motif_{i,t-1}$, to the specification, as it is possible that behavioral momentum from the previous period could influence the outcome variable seen in this period in a manner that is not captured by the lagged outcome variable alone:

$$y_{i,t} = \beta_0 + \beta_1 * motif_{i,t} + \beta_2 * used\_agent_{i,t} + \beta_3 * y_{i,t-1} + \beta_4 * motif_{i,t-1} + \beta_5 * a_i + u_{it}$$

We also propose to add an interaction term between agent usage and motif, as what we have seen earlier indicates that the direction of change of the outcome variable, if any, can depend on what motif that transaction period is associated with. The interaction term helps answer the question, "what happens when one is motif X *and* uses an agent," which provides additional granularity on top of the average motif effect or agent effect captured by the terms *motif* and *used_agent*.

$$y_{i,t} = \beta_0 + \beta_1 * motif_{i,t} + \beta_2 * used\_agent_{i,t} + \beta_3 * y_{i,t-1} + \beta_4 * motif_{i,t-1} + \beta_5 * motif_{i,t} * used\_agent_{i,t} + \beta_6 * a_i + u_{it}$$

Finally, we note the issue of **endogeneity**, which cannot really be specified in an equation as above. Endogeneity can result either from omitted variable bias, which we noted above, or from simultaneity, where the independent variables and outcome variables may influence each other. We have seen how transaction patterns differ across motifs. For example, Accumulators have many small deposits followed by large withdrawals, while Slow Drawdowns have a few large deposits followed by many small withdrawals. Since the motifs are connected to different magnitudes of the outcome variable, we cannot treat motifs as being completely exogenous.

In the following sub-sections, we explore how to address the issues of endogeneity.

Time-invariant Omitted Variables

We had intuitively described the fixed effects, a_i, that affects y_{it}, as denoting "fixed effects," reproduced below:

$$y_{i,t} = \beta_0 + \beta_1{*}motif_{i,t} + \beta_2{*}used\_agent_{i,t} + \beta_3{*}y_{i,t-1} + \beta_4{*}motif_{i,t-1} + \beta_5{*}motif_{i,t}{*}used\_agent_{i,t} + \beta_6{*}a_i + u_{it}$$

Specifically, we are looking for correlation between a_i and any one or more of the other explanatory variables (Wooldridge 2013, 474):

$Cov(x_{,it}, a_i) \neq 0$, for t = 1, 2, …, T; $x \in \{motif, used\_agent\}$

In case a_i is uncorrelated with any of the explanatory variables, (ii) would become a "random effects" model where:

$Cov(x_{i,t}, a_i) = 0$, where t = 1, 2, …, T; $x \in \{motif, used\_agent\}$

Assuming random effects assumptions hold as per (vi), we can then define a "composite error term" $v_{i,t} = a_i + u_{i,t}$ and rewrite (ii) as (Wooldridge 2013, 475):

$$y_{i,t} = \beta_0 + \beta_1{*}motif_{i,t} + \beta_2{*}used\_agent_{i,t} + \beta_3{*}y_{i,t-1} + \beta_4{*}motif_{i,t-1} + \beta_5{*}motif_{i,t}{*}used\_agent_{i,t} + v_{it}$$

Being able to use a random effects model would be a welcome result as it would address the limitation of our dataset not having any time-invariant characteristics of our savers by incorporating such fixed effects into the error term, and absolving us from having to calculate their effect directly.

We can test for whether accounting for fixed effects is necessary by using the Hausman specification test. The Hausman test "tests the null hypothesis that the coefficients estimated by the efficient random effects estimator are the same as the ones estimated by the consistent fixed effects estimator," implying that a significant p-value would indicate the presence of fixed effects (Data and Statistical Services n.d.). We run *xtreg* in stata on (ii) with options *fe* and *re*, and compare the chi-squared score and associated p-values. Table 20 show that we need to account for fixed effects across the board as p-values are highly significant.

Metric	Chi-squared Score	p-value
Avg. Balance	89,899.14	0.0000
Log (Num.Deposits)	12,721.88	0.0000
Log (Num. W/d)	27,213.61	0.0000
Log (Avg, Deposit Amt)	40,037.07	0.0000
Log (Avg, W/d Amt)	77,797.60	0.0000

Table 20. Hausman test for Fixed vs Random Effects

It thus confirms that we cannot ignore time invariant characteristics, and thus must explicitly account for them somehow.

Correlation with Lagged Values

The technical definition requires that for no autocorrelation (or serial correlation) to exist, "conditional on **X**, the errors in two different time periods are uncorrelated: $\text{Corr}(u_{it},u_{is}|\mathbf{X}) = 0$, for all $t \neq s$" (Wooldridge 2013, 341). This correlation between the errors in adjacent time periods can be written as (Wooldridge 2013, 399):

$$u_t = \rho * u_{t-1} + e_t, \text{ for } t = 1, 2, \dots, n \qquad \text{(iii)}$$

When only one adjacent time period is considered, this serial correlation model is the autoregressive model of order one, AR(1). The null hypothesis for AR(1) as specified is given below (Wooldridge 2013, 403) – i.e. no autocorrelation exists if the coefficient of u_{t-1} is zero:

$$H_0: \rho = 0$$

Since we do not know the number of lagged values of the outcome variables autocorrelation exist, we need to be able to explore a version of (iii) for additional lags. This can be done by testing for serial correlation in autoregressive models of order *q*, AR(*q*) (Wooldridge 2013, 407):

$$u_t = \rho_1 * u_{t-1} + \rho_2 * u_{t-2} + \dots + \rho_q * u_{t-q} + e_t, \text{ for } t = 1, 2, \dots, n$$

Where the null hypothesis is:

$$H_0: \rho_1 = 0, \rho_2 = 0, \dots, \rho_q = 0$$

We perform the AR() test for autocorrelation for our five outcome variables, where the null hypothesis is that there is no autocorrelation. We run this test using a user-defined Stata routine, *abar* on (ii) for up to five lags (Roodman 2006). The technique that *abar* implements was specifically designed by Arellano and Bond to be used "after estimating a dynamic model from panel data by the generalized method of moments (GMM)" (Arellano and Bond, Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations 1991). We have not discussed the Arellano-Bond approach to GMMs yet, but as we will soon see, it is our technique of choice to explore the account-level panel data on our hands. We therefore use the same test as we would in these later sections to keep the results comparable.

From the z-scores and associated p-values of the AR() test, it is clear that there is autocorrelation present for all five lags (Table 21). Our choice of five lags is somewhat arbitrary, though we feel it sufficient to illustrate pervasive autocorrelation. Note that AR(*n*) represents the AR() test for the n_{th} order.

z-score (p-value)	AR(1)	AR(2)	AR(3)	AR(4)	AR(5)
Avg. Balance	-48.32 (0.0000)	37.24 (0.0000)	33.78 (0.0000)	29.84 (0.0000)	29.79 (0.0000)
Log (Num.Deposits)	-24.65 (0.0000)	21.14 (0.0000)	18.64 (0.0000)	17.54 (0.0000)	15.98 (0.0000)
Log (Num. W/d)	-37.12 (0.0000)	35.61 (0.0000)	29.42 (0.0000)	27.53 (0.0000)	24.55 (0.0000)
Log (Avg, Deposit Amt)	-43.41 (0.0000)	31.16 (0.0000)	25.93 (0.0000)	27.80 (0.0000)	24.81 (0.0000)
Log (Avg, W/d Amt)	-49.38 (0.0000)	38.64 (0.0000)	28.75 (0.0000)	31.49 (0.0000)	27.27 (0.0000)

Table 21. Arellano-Bond test for Autocorrelation

The results of this test identify additional issues of concern – we must correct for autocorrelation. We will need to reconcile this with the fact that our sample is unbalanced, in that we have a different number of motifs for different accounts, and we can no longer hope to get away with looking at the first lag alone.

Arellano Bond GMM Estimators

Based on the work of Arellano, Bond, Bover and Blundell that utilizes General Method of Moments (GMM), we have tools available that can estimate parameters of our panel data, taking into consideration the issues outlined above, as well as others that may crop up when dealing with panel time series data (Arellano and Bond, Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations 1991) (Arellano and Bover 1995) (Blundell and Bond 1998). These GMM-based estimators can handle the following situations (Roodman 2006, 1):

1. “Small T, large N” panels, where there are few time periods and many individual units. Our dataset consists of 333,690 accounts (the “N”), while the mean and median motifs available for them are seven and ten respectively (the “T”).
2. A “linear functional relationship,” which is true for our specification.
3. One outcome variable that is dynamic, i.e. it “depends on its own past realizations.” We explained in section *The Need to Delve Further* why we think it is reasonable to expect lagged outcome variables to show up on the right hand side.
4. “Explanatory variables that are not strictly exogenous,” which we expect to be the case given a panel data with fixed effects and a lagged dependent variable. Intuitively, we can think of this applying to motifs as the shapes of the motifs are related to the number and size of deposits and withdrawals by construction.
5. “Fixed individual effects,” which we have demonstrated the presence of using the Hausman test in section *Time-invariant Omitted Variables*.
6. “Heteroskedasticity and autocorrelation within individual units’ errors, but not across them.” We demonstrated the presence of autocorrelation in our dataset in section *Correlation with Lagged Values*. We did not explore heteroskedasticity as autocorrelation within the error terms generally invalidates tests for it and can only be tested for after corrections are undertaken for autocorrelation (Wooldridge 2013, 421). Since our choice of technique itself is not invalidated by the presence of heteroskedasticity, we proceed without seeking explicit proof of its presence.

In addition, GMM estimators “do not assume that good instruments are available outside the immediate dataset”, but rather that “the only available instruments are ‘internal’ – based on the lags of the instrumented variables” (Roodman 2006, 14). We will see shortly how lagged variables are used as instruments.

Theoretical construct

While we will not detail the matrix manipulations and linear algebra constructs that allow calculation of the GMM estimators, we will explore the transformations undertaken by this technique that address the issues of our dataset before we look at the results.

To address fixed effects, GMM utilizes first-differences of the original model in (ii). We rewrite (ii) as (v) below, collapsing the regressors into a composite independent variable, *x*, for sake of concise presentation:

$$y_{i,t} = \beta_0 + \beta_1 * y_{i,t-1} + \beta_2 * x_{i,t} + \beta_3 * a_i + u_{it} \qquad \text{(v)}$$

Taking (v) for time *t-1*, and subtracting it from (v) for time *t* gives us:

$$(y_{i,t} - y_{i,t-1}) = (\beta_0 - \beta_0) + (\beta_1 * y_{i,t-1} - \beta_1 * y_{i,t-2}) + (\beta_2 * x_{i,t} - \beta_2 * x_{i,t-1}) + (\beta_3 * a_i - \beta_3 * a_i) + (u_{it} - u_{it-1})$$

Cancelling the constant terms simplifies the equation to:

$$(y_{i,t} - y_{i,t-1}) = (\beta_1 * y_{i,t-1} - \beta_1 * y_{i,t-2}) + (\beta_2 * x_{i,t} - \beta_2 * x_{i,t-1}) + (u_{it} - u_{it-1})$$

Which can them be written as:

$$\Delta y_{i,t} = \beta_1 * \Delta y_{i,t-1} + \beta_2 * \Delta x_{i,t} + \Delta u_{it} \qquad \text{(vi)}$$

Δ denotes the change from time *t-1* to *t* in (vi). This equation is called the first-difference equation. For our purposes, we find that the fixed effect, a_i, has been “differenced away” (Wooldridge 2013, 445)! This implies that as a result of this treatment, we no longer have to be concerned with potential time-invariant unobserved regressors simply because they no longer feature in the model GMM deals with. One effect of first-differencing is that we lose the very first observation, as it does not have a precursor available to difference with.

A simple first-difference transform, however, runs into issues for unbalanced datasets such as ours, where there are gaps, as there are some segments for which we may not have either balance or transaction data. In those cases, both $\Delta y_{i,t}$ and $\Delta y_{i,t+1}$ will be missing in the transformed data if $y_{i,t}$ is missing in the original. Arrelano and Bover propose a "forward orthogonal deviations" (FOD) transform, to address this, where "instead of subtracting the previous observation from the contemporaneous one, it subtracts the average of all future available observations of a variable". The FOD transform thus minimizes data loss by ensuring that all observations for each account are utilized, except the very last one (Roodman 2006, 18).

Now we turn to addressing the issue of endogeneity. First-differencing or FOD does not take care of this, as the lagged dependent variable is still endogenous due to its presence in the differenced terms. Specifically, the $y_{i,t-1}$ term in $\Delta y_{i,t-1} = (y_{i,t-1} - y_{i,t-2})$ correlates with the $u_{i,t-1}$ in $\Delta u_{i,t} = (u_{i,t} - u_{i,t-1})$. However, neither $y_{i,t-2}$ nor $\Delta y_{i,t-2}$ are related to the error term $\Delta u_{i,t}$, as long as $u_{i,t}$ are not serially correlated. This creates the opportunity to use the levels and differences of the second lag as instruments as part of a two-stage least square (2SLS) approach (Anderson and Hsiao 1982). Roodman notes that using the levels estimator is preferable as it allows for an additional time period of data, which can be significant in short panels. The efficiency of the estimators can be further improved by taking deeper lags of the dependent variable as additional instruments, but doing so within the 2SLS framework reduces the sample size as observations without lagged counterparts are dropped (Roodman 2006, 21).

Holtz-Eakin et. al. propose a solution to this by relying on the General Method of Moments framework that allows inclusion of "all valid lags of the untransformed variables as instruments, where available" (Holtz-Eakin, Newey and Rosen 1988) (Roodman 2006, 25). These "GMM-style" instruments eliminate the tradeoff between lag depth and sample depth. As part of the process, the exogenous regressors instrument themselves as "IV-style" instruments. Arellano and Bond built on this approach and introduced a two-step process to handle issues with the differenced error term, $\Delta u_{i,t}$, that can greatly distort coefficient estimates after differencing (Arellano and Bond 1991). Specifically, the two-step process ensures that “the standard covariance matrix is robust to panel-specific autocorrelation and heteroskedasticity” (Mileva 2007). These techniques together comprise of the “difference GMM estimator” for dynamic panels. Roodman confirms that for the first-differenced transform, “deeper lags of the regressors remain orthogonal to the error, and available as instruments,” and for the FOD transform, “since lagged observations do not enter the formula, they are valid as instruments” (Roodman 2006, 18).

“Difference GMM” has a shortcoming – Blundell and Bond showed that if our outcome variables are anything close to a random walk, then it “performs poorly because past levels convey little information about future changes, so that untransformed lags are weak instruments for transformed variables” (Blundell and Bond 1998) (Roodman 2006, 26). On the other hand, it is possible that for “random walk–like variables, past changes may indeed be more predictive of current levels than past levels are of current changes, so that the new instruments are more relevant,” leading Blundell and Bond to suggest instrumenting levels with differences (Roodman 2006, 27). They design a “system estimator” that consists of a dataset created by combining the transformed (either differenced or orthogonal) observations with the untransformed ones, with appropriate allowances for GMM-style and IV-style instruments (Roodman 2006, 28).

We attempted to determine if outcome variables have “unit roots” that would indicate whether segment metrics displayed a random walk or not, but the tests were inconclusive as the unbalanced nature of our dataset were not amenable to exploration using Stata commands. Anecdotally, we see accounts that have fairly steady average balances across segments through time, and we also see accounts whose balances increase or decrease on the whole. Because the “system estimator” is additive, in that prior estimators are also run, and the number of instruments remain manageable, we do not consider it to be detrimental to include in the process.

Once the "system" GMM estimators are obtained, two tests need to be run. The first is checking for autocorrelation in the "idiosyncratic disturbance term." This idiosyncratic disturbance term is composed of two orthogonal components – the fixed effects, which the estimators should eliminate, and the idiosyncratic shocks, which are captured in the residual error. The implication for finding autocorrelation of order 1 is that we would need to use lags 3 and deeper as instruments. (Roodman 2006, 14, 32) We can generalize this to saying that for autocorrelation of order k, we would need to start with lags $k + 2$ and deeper as instruments. Arellano and Bond offer a test to check for autocorrelation in the idiosyncratic disturbance term, which we will refer to in our analysis.

The second involves the "Sargan-Hansen test for over-identifying restrictions." The joint null hypothesis for the Sargan-Hansen test is that "the instruments are valid instruments, i.e., uncorrelated with the error term, and that the excluded instruments are correctly excluded from the estimated equation" (StataCorp 2009). There is some difference of opinion as to what this implies. Mileva interprets this null hypothesis to mean the same as "the instruments as a group are exogenous" (Mileva 2007). Parente and Silva caution that the Sargan-Hansen test is *not* a test for exogeneity, and that "the validity of the overidentifying restrictions provides little information on the ability of the instruments to identify the parameter of interest," "the interpretation of the outcome of a test for overidentifying restrictions does not depend on the presence of enough valid instruments," and "it is more appropriate to interpret tests for overidentifying restrictions as checks for whether or not all the instruments identify the same vector of parameters" (Parente and Silva 2011). Angrist and Pischke go so far as to say that such overidentification testing "is out of the window in a fully heterogenous world" (Angrist and Pischke 2008, 166). We will restrict our interpretation of the Sargan-Hansen test as simply to check that the instruments are not correlated with the error term. Note that the Sargan-Hansen test still assumes that one of the instruments is exogenous – an assumption it cannot test for too. In our case, that instrument is agent usage.

The test is also weakened by the presence of many instruments. It is not clear how many is too many, though one rule of thumb mentioned by Roodman is that the number of instruments should not exceed the number of accounts in our panel. The number of instruments is quadratic to the number of periods, T, and the size of one of the key moment matrices is quadratic to the number of instruments, and therefore *quartic* to T. A large instrument count can also overfit endogenous variables. (Roodman 2006, 12, 13). On a practical note, we can (and have) run out of memory as the process attempts to fit ballooning matrices of moment conditions. We therefore make a point of using as few lags as possible as instruments to obtain a reliable Hansen statistic, reduce the danger of overfitting, and actually complete the runs on available hardware.

Methodology

Before we derive AB GMM estimators for our dataset[2], we prune it in two ways to fit it all into memory. We remove all accounts which have either one or two periods only. Because we can only instrument starting with the second lag, these 15,993 accounts would never feature in the GMM estimators anyway. This leaves us with 23,987 accounts.

---

[2] We use Roodman's *xtabond2* implementation of the GMM estimators instead of Stata's built in *xtabond* command. *xtabond2* can do everything *xtabond* can, and has a few additional advantages. It can apply a "Windmeijer finite-sample correction to the standard errors in the two-step estimation, without which those standard errors tend to be severely downward biased," allows for the FOD transformation that "preserves sample size in panels with gaps," and specify lags to include for GMM-style instruments with more granularity (Roodman 2006, 1) (Baum 2013, 20).

These accounts have 312,792 segments between them, with the most prolific account having 175 segments! We noted earlier how some of the moment matrices will grow quartic to the number of segments, turning out to be the primary cause of running out of memory. To manage this issue, we discard all segments beyond the 34^{th}, for all accounts. Thirty-five and beyond represent the top 5% of segments, meaning that we still have 95% of segments to run system GMM on. Note that none of the 23,987 accounts are discarded, only segments thereof when they are not in the first 35.

With this sample, we run Arellano-Bond "system" GMM two-step estimators, with FOD transforms and Windmeijer's finite-sample correction. We consider instruments created through levels of motifs and agent usage only, and not differences, as "differences" between designations of a categorical variable does not make sense – we can't really subtract the motif, Accumulator, from the motif, Sustained Balance, for example.

It is worth pointing out that we treat out *used_agent* explanatory variable as exogenous. We feel comfortable doing so we find that there is no evidence in any of our runs that it is correlated with the error term.

We test four specifications: with agent usage but no motifs (1), with motifs but no agent usage (2), with motifs and agent usage but no interaction between the two (3), and finally, with motifs, agents, and interaction between the two (4) below. We are really only interested in (4) as that captures the interaction between motifs and agent usage. The results for (1) and (2) are presented for average balances only for illustration purposes, but second order autocorrelation is quite present in both for all five outcome variables, rendering the results unusable. We continue to present (3) along with (4) as the changes coefficients assisted us in understanding the effect associated with agents better in (4); (3) does not provide additional pedagogic value otherwise.

Once we have our results, we check the Hansen test to see if our instruments continue to be correlated with the error term. If they are, we inspect the histogram for residuals, a standardized normal probability plot (P-P) and a quintile plot of errors to normal distribution (Q-Q). The standardized normal probability plot (P-P) is sensitive to non-normality in the middle range of the data, while the Q-Q plot plots the quintiles of a variable against the quintiles of a normal distribution and is sensitive to non-normality near the tails (UCLA: Statistical Consulting Group n.d.).

If we do find any of the explanatory variables to be correlated with the error term, we diagnose whether it results in an upward bias or a downward bias. We leave the extended discussion till we can illustrate the situation with actual results from average balances, but we note here that for lagged outcome variables, the bias is always upward, i.e. the coefficient overestimates the effect associated with the lagged outcome variable.

Complete details of the AB-GMM estimations are presented in *Appendix B. Detailed Results of AB-GMM Runs*. The next four sections draw from these results.

Average Motif Balance

We present the system GMM estimators with average balance as outcome variable below, with agent usage but no motifs (1), with motifs but no agent usage (2), with motifs and agent usage but no interaction between the two (3), and finally, with motifs, agents, and interaction between the two (4).

Explanatory Variables	(1)	(2)	(3)	(4)
Lagged log(balance)	0.019 *** (0.004)	0.0739 *** (0.0188)	– 0.108*** (0.030)	– 0.101 *** (0.031)
Agent Used	0.148 *** (0.016)		– 0.021 (0.035)	– 0.519 (0.765)
M2: Dump-and-Pull		– 1.734 *** (0.241)	– 1.798 *** (0.385)	– 2.003 *** (0.505)
M3: Fast Drawdown		– 1.092 *** (0.001)	– 1.676 *** (0.512)	– 2.065 *** (0.622)
M4: Accumulator		– 2.713 *** (0.300)	– 3.624 *** (0.531)	– 3.768 *** (0.680)
M5: Slow Drawdown		– 0.197 (0.277)	– 0.597 (0.540)	– 0.707 (0.675)
Lagged M2: Dump-and-Pull		– 0.159 (0.216)	– 1.864 *** (0.344)	– 1.814 *** (0.355)
Lagged M3: Fast Drawdown		– 0.835 *** (0.277)	– 2.495 *** (0.406)	– 2.455 *** (0.426)
Lagged M4: Accumulator		– 0.355 (0.381)	– 1.118 (0.715)	– 1.013 (0.737)
Lagged M5: Slow Drawdown		– 1.399 *** (0.270)	– 1.783*** (0.517)	– 1.793 *** (0.527)
Agent Used × M2				0.547 (0.991)
Agent Used × M3				1.454 (1.294)
Agent Used × M4				0.463 (1.349)
Agent Used × M5				0.315 (1.064)
Intercept	7.908 *** (0.034)	9.071 *** (0.347)	11.882 *** (0.537)	11.972 *** (0.612)
AR(1)	0.000	0.000	0.000	0.000
AR(2)	0.000	0.000	0.558	0.671
Hansen test	0.000	0.000	0.020	0.017
Wald χ^2	0.0000	0.0000	0.0000	0.0000
F-test for Motif Joint Sig.	-	0.0000	0.0000	0.0000
F-test for Joint Sig., Lagged Motifs	-	0.0000	0.0000	0.0000
F-test for Joint Sig., Interaction	-	-	-	0.8573
No. of observations	273,820	228,481	228,481	228,481
No. of accounts (groups)	23,897	23,097	23,097	23,097
No. of instruments	596	723	196	196

Standard errors are in parentheses. *, ** and *** represent significance at 10%, 5% and 1% respectively.
AR(1) – 1st order autocorrelation test; AR(2) – 2nd order autocorrelation test
(1) and (2) use all available lags as instruments; (3) and (4) use second lags only.

Table 22. System GMM Estimators for Average Balance

Figure 26 provides a more visual representation of the coefficients for (4), making it easier to compare magnitude of impact.

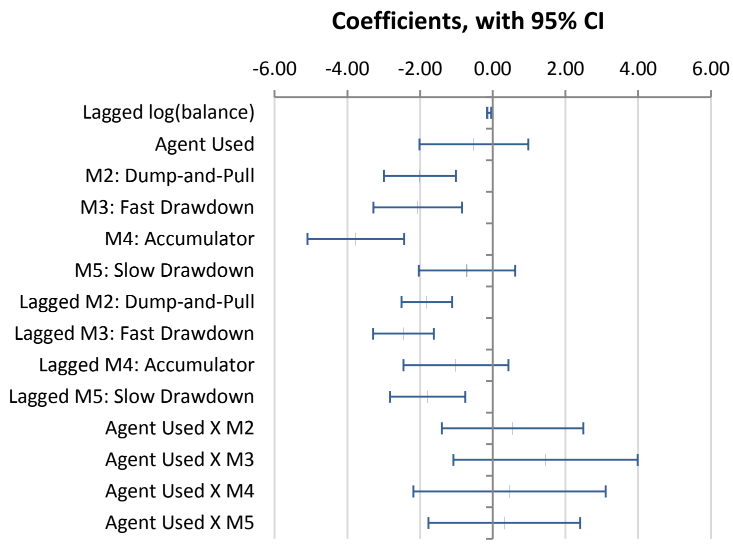

Figure 26. Coefficients with 95% CI for Average Balance

Second-order aurocorrelation for average motif balances is mitigated through GMM estimators in (3) and (4) only, and not in (1) and (2). The Hansen test is weak, suggesting that our instruments continue to be correlated with the error term. The number of instruments, 196, is much smaller than the number of accounts, at 23,097, implying that we do not have a "too many instruments" problem. Let us look at each component of our model, and then explore the implications of the weak Hansen test.

- log(*balance*$_{i,t-1}$) : An approximately 10% reduction in balance can be expected from the previous period's balance. This implies that balances go down over time.
- *used_agent*$_{i,t}$: Agent usage is not significantly associated the average balance levels in a segment. This implies that there is no correlation between the balance of a period and whether an agent was used in that period.
- *motif*$_{i,t}$: Three of the four motif coefficients are statistically significant, and the F-test for joint significance, which has the null hypothesis that all coefficients on motifs are statistically indistinguishable from zero, is significant at the 1% level. This suggests that

motifs have statistically significantly different balance levels from each other. Because Sustained Balances have the highest average balances and it is the omitted motif in the regression run, all the other coefficients come out to be negative, as they are all reported compared to Sustained Balance motifs. The average balances of Slow Drawdown motifs are not statistically significantly different than those of Sustained Balances.

- *motif*$_{i,t-1}$: Three of the four lagged motif coefficients are statistically significant, and the F-test for joint significance for lagged motifs is also significant, at the 1% level. This suggests that what motif an account was in the previous segment is correlated with the balance levels of the current segment. We can interpret the motif and lagged motif coefficients together as follows – if an account had a Sustained Balance motif this period and the previous one, its logged balance value would be 11.972. If it was a Dump-and-Pull motif for both the current period and the previous one, we can expect a balance of (11.972 – 2.003 – 1.814) or 8.155. We can construct a 4 x 4 matrix which provides expected logged balance amounts compared to Sustained Balance segments. We refrain from exploring this in more detail as it is sufficient for our purposes to demonstrate that motifs, current and previous, *are* significantly correlated with current balance levels.
- *motif*$_{i,t}$**used_agent* : None of the interaction terms between agents and motifs is statistically significant, and the F-test for joint significance has a p-value of 0.8573, suggesting the coefficients are highly indistinguishable from zero. This implies that we cannot discern any significant relationship between the use of agents within the context of a particular motif and the average balance in that motif, given what we already know about what motif the account is and whether it has used an agent in this segment.

This is confirmed in Table 23 below, where we present the pairwise significant tests for interaction terms. The coefficients of *motif*$_{i,t}$**used_agent* already tells us if each of the interaction of any motif is different when compared to that of Sustained Balances. Pairwise comparisons extend this check to all six pairs possible from the remaining four motifs. While we don't extend this check to the absolute values of the coefficients, comparing their relative differences to Sustained Balances still allows us to determine if they have a differentiated effect associated with them. Only the upper quadrant is presented as the results are symmetrical across the diagonal, and therefore redundant.

Prob($\beta_{5,X} = \beta_{5,Y}$)			
Y	X		
	Fast Drawdown	Accumulator	Slow Drawdown
Dump-and-Pull	0.414	0.942	0.780
Fast Drawdown		0.514	0.362
Accumulator			0.908
Pairwise comparisons are being made across the coefficients of motifs X and Y *, ** and *** represent significance at 10%, 5% and 1% respectively.			

Table 23. Pairwise Comparison of Interaction Coefficients, Avg. Balance

We thus find that current balance levels are significantly correlated with previous balance levels, but not with agent usage during the current segment. Current and previous motifs show differentiated behavior amongst themselves. Additional distinction is not provided through motif-agent interactions.

As we have seen though, the Hansen test tells us that our instruments are correlated with the error term. Let us explore this issue closely as it has the potential to impact the validity of our entire specification.

The error terms seem to be fairly normally distributed. The histogram of errors with a normal density plot superimposed demonstrates a slight skew Figure 27 (a). The P-P plot is quite linear, suggesting normality in the mid-range of the errors (Figure 27 (b)), while there is some disturbance in the lower tail evident in the Q-Q plot (Figure 27 (c)). The residuals can therefore be considered to be approximately normal, with minor deviations.

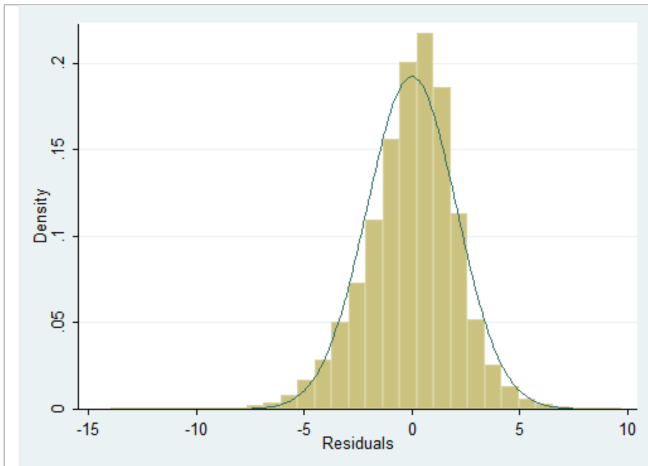

(a) Histogram for residuals, with scaled normal density plot

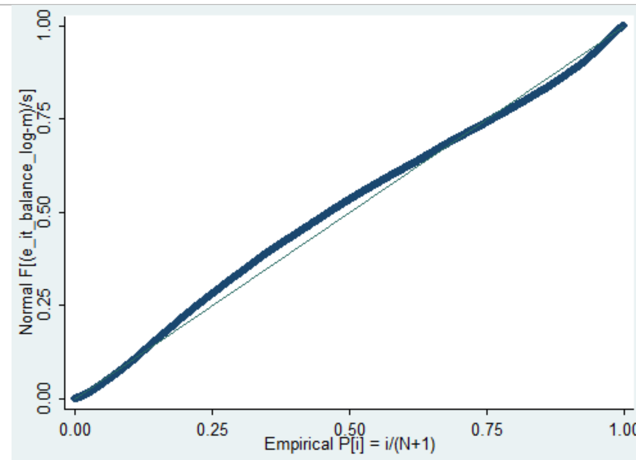

(b) Standardized normal probability plot (P-P)

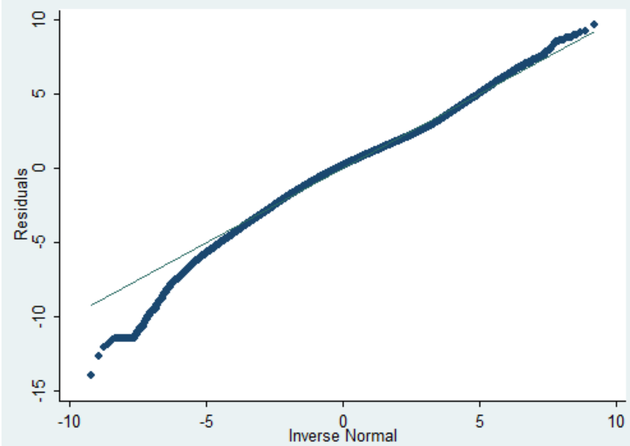

(c) Quintile plot of errors to normal distribution (Q-Q)

Figure 27. Error Plots, Average Balance

Since balance values are considered to be endogenous and we expect correlation with the first lagged values, we test residuals against the second lagged value of balances. The upward slope is quite evident in Figure 30 (a). Regressing the second lagged logged balance to residuals gives following, with a very weak R^2 of 0.0946:

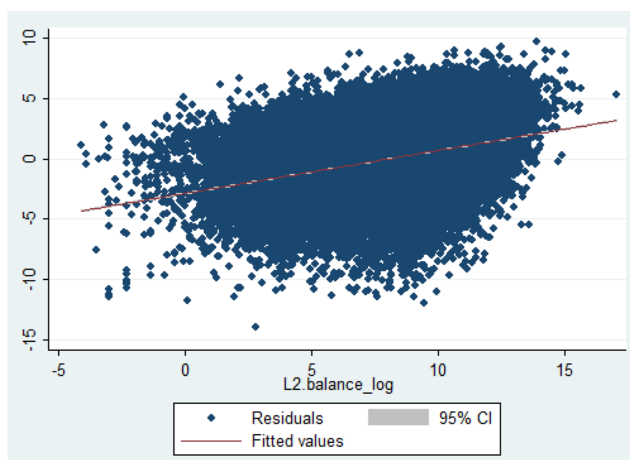

Figure 28. Residuals plotted against second lag of log(balance)

$$e_{it} = -2.650 + 0.329 * \log(balance_{i,t-2})$$

Both the constant and coefficient are statistically significant at the 1% level, and the coefficient on $balance_{i,t-2}$ confirms the upward slope. Adding additional lags does not instrument this correlation away. We will come back to this shortly, but first, let us determine if there are issues between the other explanatory variables and the error term.

Since agent usage is considered to be exogenous and instrumented with itself alone, we can compare the residuals with the current agent usage. It seems that the errors are centered on zero from the boxplot in Figure 29 (a), and have a normal distribution in Figure 29 (b). There therefore does not seem be an issue with errors as far as agent usage is concerned.

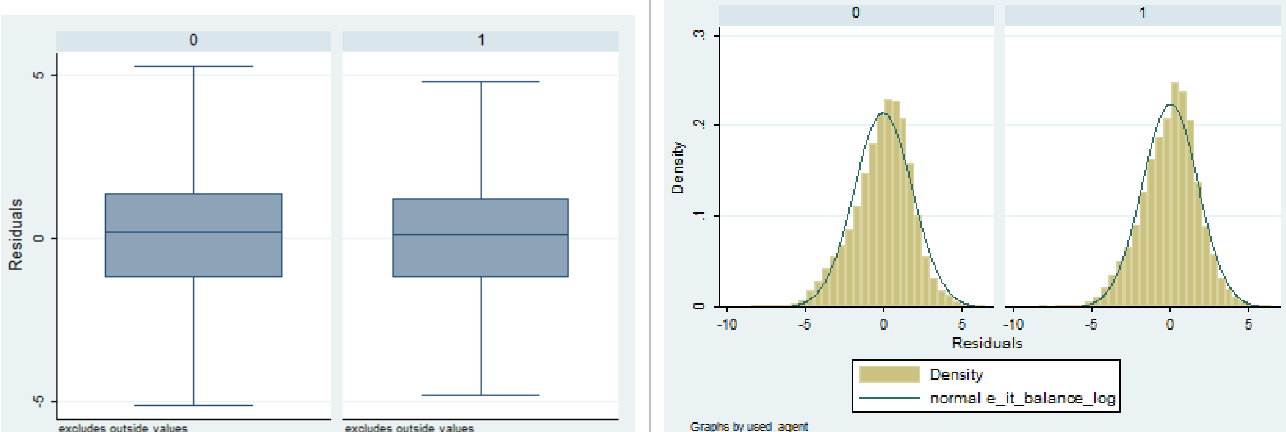

(a) Boxplot (b) Histogram with normal density plot

Figure 29. Residual Plots by Agent Usage

Note that we do not present the F-test for mean errors by agent usage, or any other covariate, as further proof of zero-ness of mean because our high sample size of 228,481 gives these tests extraordinary power and arrives at extremely narrow confidence intervals as a result, suggesting effects that are statistically significant but also very small in magnitude. Thus, e_{it} has a 5^{th} to 95^{th}

percentile range of [-3.668, 3.092] with a mean of -0.011 in actual data, while the 95% confidence interval reported is [-0.024, -0.004].

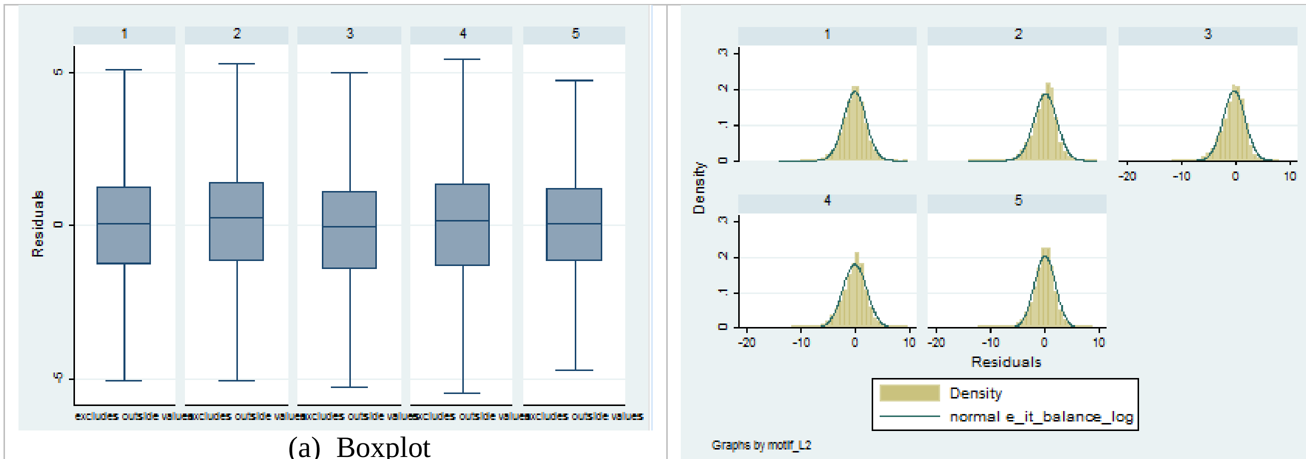

(a) Boxplot

(b) Histogram with normal density plot

Figure 30. Residual Plots against second Lagged Motif

Motifs are suspected to be endogenous, and therefore residuals are tested against motifs from two lags prior. We can see slight deviations in the boxplots for residuals for motifs 2 (Dump-and-pull) and 4 (Accumulators) over second lagged motifs in Figure 30 (a). The residuals look normally distributed Figure 30 (b). Therefore there does not seem to be any major issues with residual distribution as far as motifs are concerned.

This leaves us with the predicament that the second lagged balance is correlated with the error term, though not any of the other covariates. A linear relationship such as seen in Figure 28 is often indicative of an omitted variable bias. We have already accounted for time-invariant omitted variables through fixed-effects treatment, which suggests that this particular omitted variable is time-*variant* – its effect on balance changes over time. Can we predict how this impacts the coefficients on our specification?

Given an outcome variable Y, an included independent variable X_i and an omitted independent variable X_o, bias will appear in the following manner on X_i's coefficients, β_i [adapted from (Wooldridge 2013, 86)]:

	Positive Correlation, X_i and X_o	Negative Correlation, X_i and X_o
Positive Correlation, Y and X_o	β_i is overestimated	β_i is underestimated
Negative Correlation, Y and X_o	β_i is underestimated	β_i is overestimated

Table 24. Omitted Variable Bias Reference Table

But X_i is, in our case, lagged Y, as $Y = balance_{i,t}$, and $X_i = balance_{i,t-2}$. If X_i is correlated in a certain direction to X_o, so must Y when it manifests itself as such two periods later. Therefore either both are positively correlated to X_o, or negatively – but it is not possible for one to be positively correlated and the other negatively. This means the viable quadrants are the upper left or the lower right ones, which in turn means that β_i can only be overestimated in our case.

Based on the regression we ran on e_{it} over $balance_{i,t-2}$, we found above that X_i had a positive correlation with the error term, implying that it also has a positive correlation to X_o. We can therefore conclude that the appropriate quadrant is the top left.

Even though we have ascertained that β_i is overestimated, we do not explicitly see the coefficient of the instruments in our specification, and therefore must conclude this issue by noting that there is an upward bias in the instrumented variable created by lagged balances, and adjust its effect mentally towards zero. To the extent that balances are serially correlated and the same bias transports over to the first lag, we can also mentally adjust the coefficients for the lagged balance towards zero.

We do not have many guesses on what X_o could possibly be. Demographic characteristics such as education and age seem to be unlikely candidates, as we only have two years of data, and it is unlikely that changes attributable to these factors would express themselves in such a short timeframe. Ability to use and confidence in mobile money, greater availability of banking agents and a greater network effect as friends and family members sign up are more likely factors that could be positively correlated with balances. Unfortunately, our dataset does not allow us to explore the effect of such covariates.

Number and Average Amount of Deposits

It is not possible to remove persistent autocorrelation while exploring agents or motifs alone, as per specifications noted as (1) and (2) at the beginning of our discussion on average account balances. We will therefore be ignoring the two that were "agents only" and "motifs only," and discuss the specifications for agents and motifs (3), and agents, motifs and their interaction (4) only.

Deposits	**Number**		**Average Amount**	
Explanatory Variables	(3)	(4)	(3)	(4)
Lagged log(deposit metric)	0.100 *** (0.015)	0.106 *** (0.016)	0.112 *** (0.025)	0.099 *** (0.025)
Agent Used	– 0.039 * (0.022)	1.170 *** (0.448)	– 0.273 *** (0.029)	– 0.642 (0.524)
M2: Dump-and-Pull	– 0.435 * (0.237)	0.210 (0.322)	0.727 ** (0.318)	0.990 ** (0.389)
M3: Fast Drawdown	– 1.346 *** (0.287)	– 0.917 ** (0.360)	0.212 (0.371)	0.525 (0.463)
M4: Accumulator	– 2.117 *** (0.329)	– 2.143 *** (0.396)	2.957 *** (0.362)	2.511 *** (0.452)
M5: Slow Drawdown	– 0.986 *** (0.311)	– 0.390 (0.383)	1.020 ** (0.398)	0.409 (0.454)
Lagged M2: Dump-and-Pull	– 0.875 *** (0.244)	– 0.755 *** (0.252)	1.020 *** (0.269)	0.934 *** (0.273)
Lagged M3: Fast Drawdown	– 1.053 *** (0.299)	– 0.860 *** (0.311)	0.684 ** (0.319)	0.532 (0.326)
Lagged M4: Accumulator	– 0.457 (0.374)	– 0.332 (0.390)	1.341 *** (0.418)	1.135 *** (0.431)
Lagged M5: Slow Drawdown	– 0.557 * (0.312)	– 0.367 (0.261)	0.301 (0.367)	0.269 (0.371)
Agent Used × M2		– 1.697 *** (0.571)		– 0.437 (0.689)

Deposits	Number		Average Amount	
Agent Used × M3		– 1.534 ** (0.741)		0.639 (0.937)
Agent Used × M4		– 0.205 (0.767)		0.354 (0.920)
Agent Used × M5		– 1.554 ** (0.642)		2.126 *** (0.743)
Intercept	2.775 *** (0.281)	2.179 *** (0.347)	6.419 *** (0.376)	6.650 *** (0.407)
AR(1)	0.000	0.000	0.000	0.000
AR(2)	0.430	0.857	0.187	0.277
Hansen test	0.001	0.016	0.000	0.000
Wald χ^2	0.0000	0.0000	0.0000	0.0000
F-test for Joint Sig., Motifs	0.0000	0.0000	0.0000	0.0000
F-test for Joint Sig., Lagged Motifs	0.0000	0.0002	0.0000	0.0034
F-test for Joint Sig., Interaction	0.0001	0.0083	0.0006	0.0007
No. of observations	176,132	176,132	176,132	176,132
No. of accounts (groups)	21,610	21,610	21,610	21,610
No. of instruments	196	196	196	196

Standard errors are in parentheses. *, ** and *** represent significance at 10%, 5% and 1% respectively.
AR(1) – 1st order autocorrelation test; AR(2) – 2nd order autocorrelation test
(1) and (2) use all available lags as instruments; (3) and (4) use second lags only.

Table 25. System GMM Estimators for Frequency and Amount of Deposits

First, we look at the effect of motifs and agent usage on the number of deposits. . Figure 31 provides a more visual representation of the coefficients for (4), making it easier to compare magnitude of impact.

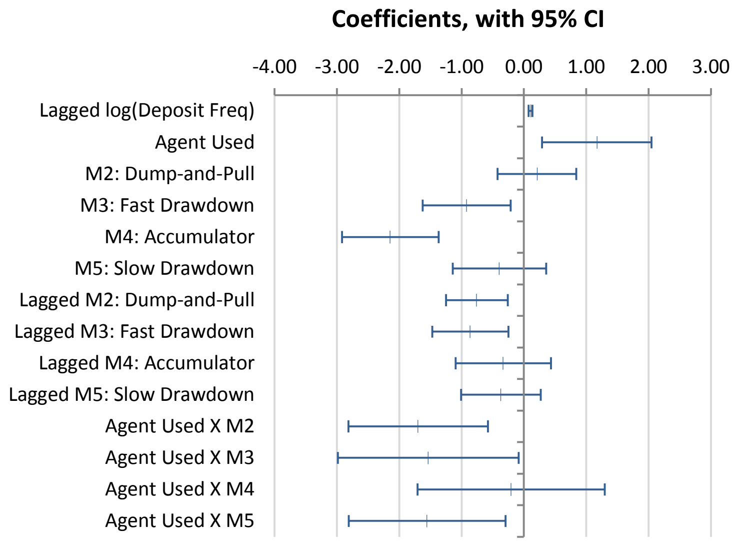

Figure 31. Coefficients with 95% CI for Deposit Frequency

Second-order aurocorrelation for number of deposits is mitigated through GMM estimators in both (3) and (4). The Hansen test is weak, suggesting that our instruments continue to be correlated with the error term. The number of instruments, 196, is much smaller than the number of accounts, at 21,610, implying that we do not have a “too many instruments” problem. Let us look at each component of our model.

- log(*number of deposits*$_{i,t-1}$) : An approximately 10% increase in the number of deposits can be expected from the previous period’s figures. This implies that the number of deposits increases over time.
- *used_agent*$_{i,t}$: Specification (3) suggests that the overall decrease in frequency is about 4% when an agent is used. The coefficient of (4) tells us that Sustained Balance segments show an increase of deposit frequency on agent usage. The other motifs cannot be looked at independently of the interaction term; discussion is deferred to the appropriate section below.

- *$motif_{i,t}$* : Two of the four motif coefficients are statistically significant, and the F-test for joint significance, which has the null hypothesis that all coefficients on motifs are statistically indistinguishable from zero, is significant at the 1% level. This suggests that at least some of the motifs have statistically significantly different deposit frequencies. We find that Fast Drawdowns and Accumulator accounts have lower deposit frequencies than Sustained Balances, while those of Dump-and-Pulls and Slow Drawdowns are indistinguishable from it.
- *$motif_{i,t-1}$* : Two of the four lagged motif coefficients are statistically significant for lagged motif too. The F-test for joint significance for these lagged motifs is also significant, at the 1% level. In this case, the Dump-and-Pull and Fast Drawdown accounts have significantly different deposit frequencies. The coefficients of *$motif_{i,t}$* and *$motif_{i,t-1}$* are best considered in tandem. Thus, assuming no agent usage, the number of deposits expected for accounts that were Sustained Balance both in the current and previous period is $e^{2.179}$, or about 8.84 per 30 days, while that for accounts which were Fast Drawdowns in both is $e^{(2.179-0.917-0.860)}$, or about 1.49 per 30 days. The relevant conclusion in this case is that current and lagged motifs are significant predictors of deposit frequencies.
- *$motif_{i,t}$*used_agent* : Three of the four interaction terms are statistically significant, as confirmed by the F-test for joint significance. The interpretation of these coefficients must also follow an additive approach, where the coefficients of the interaction term are considered jointly with the coefficient of the stand-alone agent usage. Thus, if one were a Slow Drawdown account, the log of deposit frequencies would increase by 1.170 by virtue of it having used an agent, but then decrease by -1.554, because it happens to be a "motif 5" that used the agent. However, that is not all – we must also consider what the current and lagged motifs are to arrive at figures for the remaining four motifs and not just the base Sustained Balance accounts. The overall conclusions to draw from the interaction term coefficients is that Accumulators and Sustained Balance accounts

increase deposit frequencies by similar amounts, while that of the other three are lower compared to those of Sustained Balances.

Table 26 tells us that the in interaction effect is particularly different between Accumulators and Dump-and-Pulls, and Accumulators and Slow Drawdowns. This is consistent with the observation that interaction with Accumulators is not statistically different compared to Sustained Balances, but Sustained Balances is different compared to both Dump-and-Pulls and Accumulators. Accumulators and Fast Drawdowns have an almost statistically different interaction effect, with a p-value of 0.108. The effect of Fast Drawdowns is indistinguishable from those of Slow Drawdowns, and Dump-and-Pulls.

Prob($\beta_{5,X} = \beta_{5,Y}$)			
Y	X		
	Fast Drawdown	Accumulator	Slow Drawdown
Dump-and-Pull	0.793	0.017 **	0.774
Fast Drawdown		0.108	0.976
Accumulator			0.048 **
Pairwise comparisons are being made across the coefficients of motifs X and Y *, ** and *** represent significance at 10%, 5% and 1% respectively.			

Table 26. Pairwise Comparison of Interaction Coefficients, Deposit Frequency

We thus find that deposit frequencies are significantly correlated with previous frequencies, and with agent usage during the current segment. Current and previous motifs show differentiated behavior amongst themselves, as well as additional distinction through motif-agent interactions.

The Hansen test tells us that our instruments are correlated with the error term. Let us explore this issue closely as it has the potential to impact the validity of our entire specification. The error terms seem to be fairly normally distributed. The histogram of errors with a normal density plot superimposed demonstrates a slight skew Figure 32 (a). The P-P plot is quite linear, suggesting normality in the mid-range of the errors (Figure 32 (b)). There is some disturbance in both the upper and lower tails evident in the Q-Q plot (Figure 32 (c)), though it seems to be fairly minimal. The residuals can therefore be considered to be approximately normal, with minor deviations.

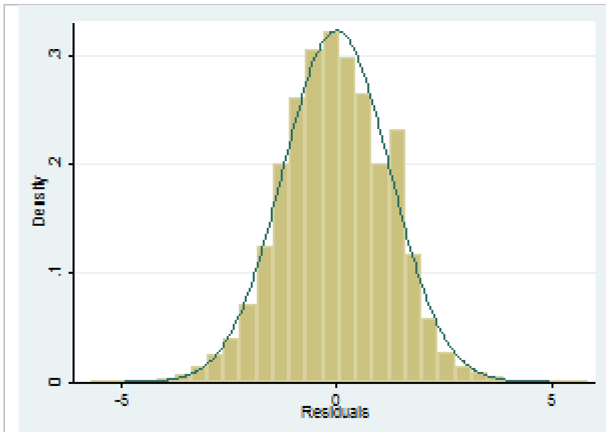

(a) Histogram for residuals, with scaled normal density plot

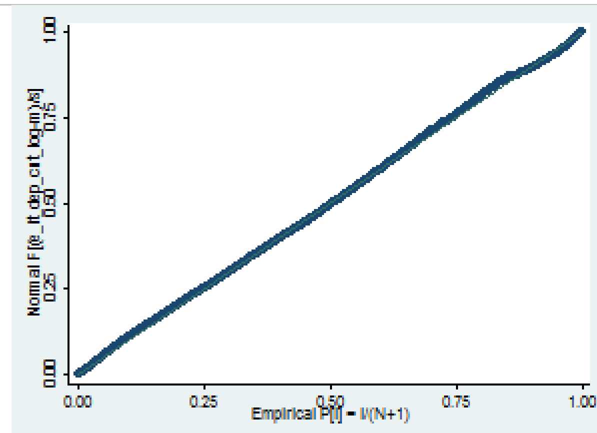

(b) Standardized normal probability plot (P-P)

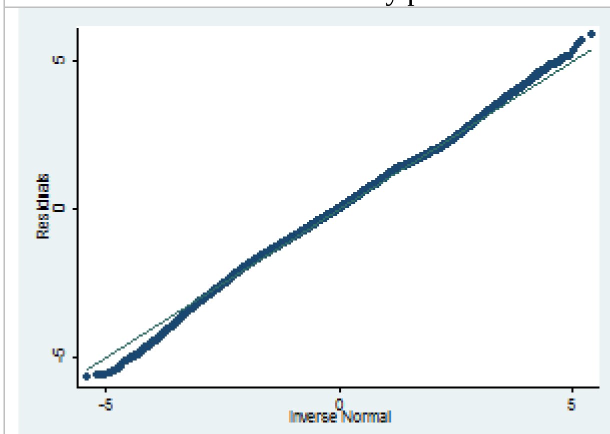

(c) Quintile plot of errors to normal distribution (Q-Q)

Figure 32. Error Plots, Deposit Frequency

Since number of deposits is considered to be endogenous and we expect correlation with the first lagged values, we test residuals against the second lagged value of deposit frequencies. The upward slope is quite evident in Figure 33. Regressing the second lagged logged deposit frequency to residuals gives the following, with a very weak R^2 of 0.0898:

$$e_{it} = -0.366 + 0.298 * \log(num\_deposits_{i,t-2})$$

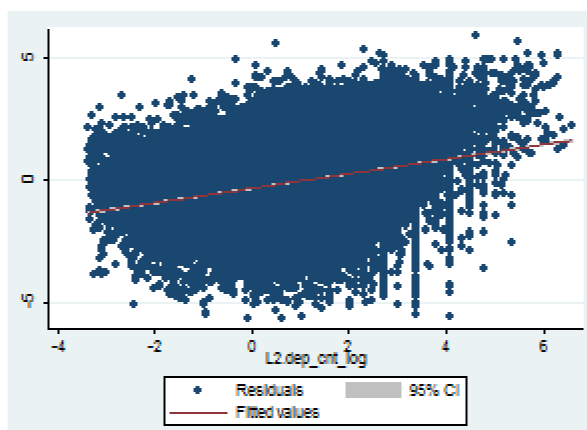

Figure 33. Residuals plotted against second lag of log(Number of Deposits)

Both the constant and coefficient are statistically significant at the 1% level, and the coefficient on *num_deposits*$_{i,t-2}$ confirms the upward slope. Adding additional lags does not instrument this correlation away. This leaves us with the predicament that the second lagged deposit frequency is correlated with the error term.

As we noted for balances, a linear relationship such as seen in Figure 33 is often indicative of an omitted variable bias, and since we have already accounted for time-invariant omitted variables through fixed-effects treatment, we must conclude that this particular omitted variable is time-*variant* – its effect on number of deposits changes over time. Relying on Table 24 and using the same logic, we can make the case that the effect of the instrumented variable in the form of the second lag of deposit frequencies will be overestimated, and must be mentally adjusted towards zero. To the extent that deposit frequencies are serially correlated and the same bias transports over to the first lag, we can also mentally adjust the coefficients for the lagged deposit frequency towards zero. Our guesses for what this omitted variable could be are no different than what we discussed for balances, and are not rehashed.

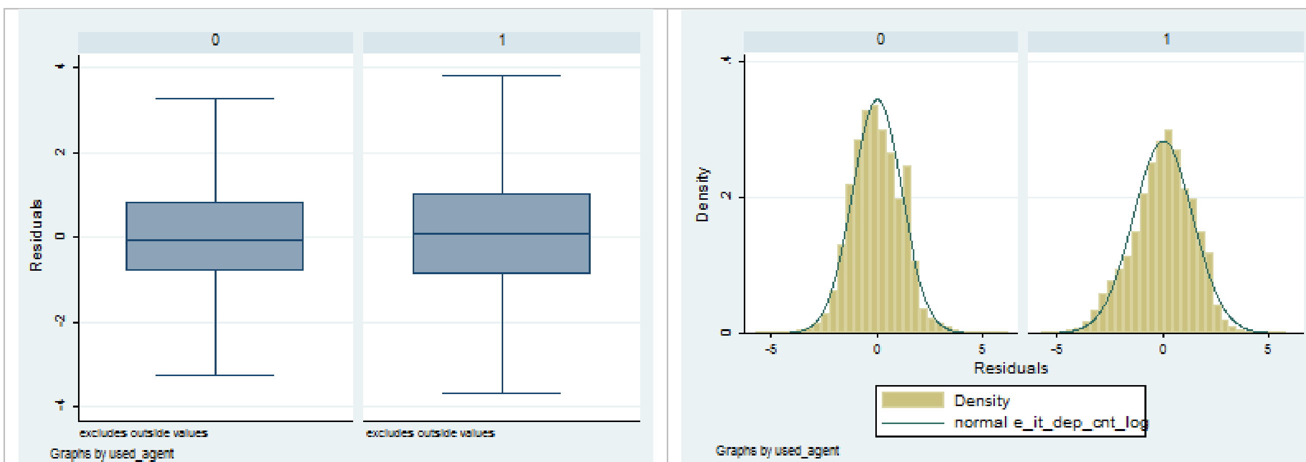

(a) Boxplot (b) Histogram with normal density plot

Figure 34. Residual Plots by Agent Usage

Since agent usage is considered to be exogenous and instrumented with itself alone, we can compare the residuals with the current agent usage. It seems that the errors are centered on zero from

the boxplot in Figure 34 (a), and have a normal distribution in Figure 34 (b). There therefore does not seem be an issue with errors as far as agent usage is concerned.

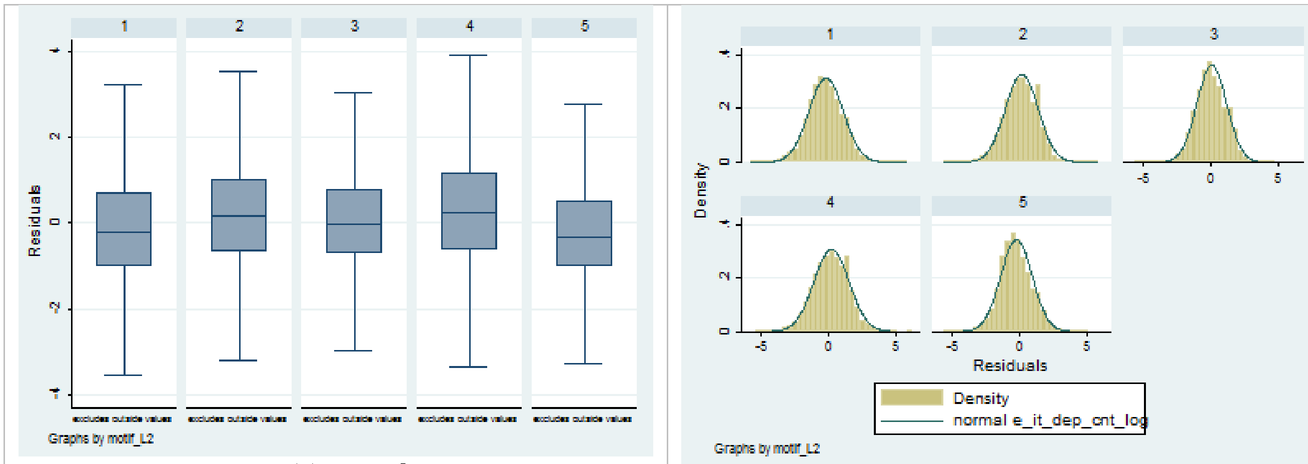

(a) Boxplot (b) Histogram with normal density plot

Figure 35. Residual Plots against second Lagged Motif

Motifs are suspected to be endogenous, and therefore residuals are tested against motifs from two lags prior. We can see some deviations in the boxplots for residuals for motifs 1 (Sustained Balance), 4 (Accumulators) and 5 (Slow Drawdown) over second lagged motifs in Figure 35 (a). The residuals look normally distributed Figure 30 (b). Therefore there seems to be minor issues with residual distribution as far as motifs are concerned.

We can extend the argument we presented earlier for the lagged outcome variable to motifs too, in that the correlation of both the outcome variable and the included independent variable must be in the same direction with the omitted independent variable as they are but lagged versions of each other, with the current values becoming the lagged values in due course. Since the correlation must be either positive for both, or negative for both, the estimate for the lagged motif instruments are overestimated, and must be corrected downward, towards zero. In so far as lagged motifs of the current run are current motifs in some other run, we can extrapolate these results to say that we must also downward correct the coefficients of $motif_t$, and $motif_{t-1}$. In other words, the decreases in deposit frequencies are not as large as they seem for motifs that sport such a decrease.

Next, we look at the effect of motifs and agent usage on the average size of deposits. Figure 36 provides a more visual representation of the coefficients for (4), making it easier to compare magnitude of impact.

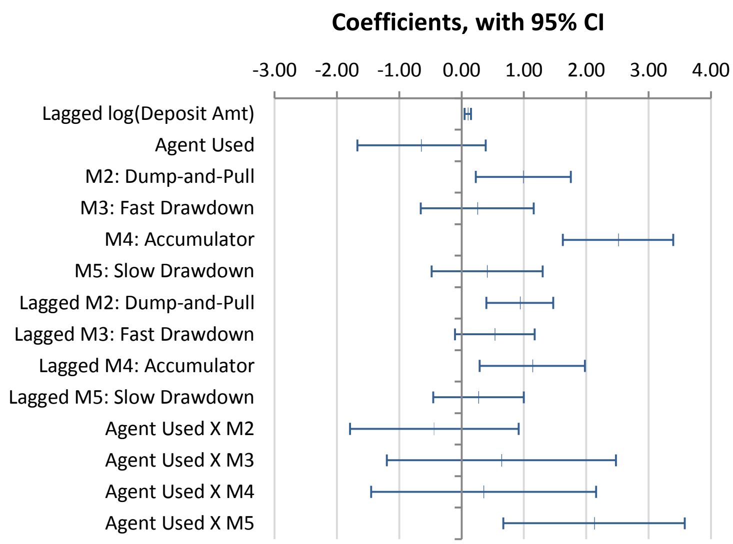

Figure 36. Coefficients with 95% CI for Avg Deposit Amt

Second-order aurocorrelation for average deposit amounts is mitigated through GMM estimators in both (3) and (4). The Hansen test is weak, suggesting that our instruments continue to be correlated with the error term. The number of instruments and accounts are the same as for deposit frequencies, implying that we do not have a "too many instruments" problem. Let us look at each component of our model.

- log($avg\_deposit_{i,t-1}$) : An approximately 10% increase in the size of deposits can be expected from the previous period's figures. This implies that the average amount of deposits increases over time.

- *used_agent*$_{i,t}$: Specification (3) suggests that the overall decrease in average amount of deposit when an agent is used in that segment, decreasing from about $e^{6.419}$ or KES 613 ($6.81) to about $e^{(6.419-0.273)}$ or KES 466 ($5.19). The coefficient of (4) is not statistically significant for Sustained Balances. The rest cannot be looked at independently of the interaction term; discussion is deferred to the appropriate section below.
- *motif*$_{i,t}$: Two of the four motif coefficients are statistically significant, and the F-test for joint significance, which has the null hypothesis that all coefficients on motifs are statistically indistinguishable from zero, is significant at the 1% level. This suggests that at least some of the motifs have statistically significantly different average deposit amounts. We find that Accumulators and Dump-and-Pull accounts have higher average deposit amounts than Sustained Balances, while those of Fast and Slow Drawdowns are indistinguishable from it.
- *motif*$_{i,t-1}$: The same two of the four lagged motif coefficients are statistically significant for lagged motif as they were with contemporary motifs. The F-test for joint significance for these lagged motifs is also significant, at the 1% level. This suggests that what motif an account was in the previous segment is correlated with higher deposit amounts in the current segment if they were Accumulators and Dump-and-Pull accounts.
- *motif*$_{i,t}$**used_agent* : The interpretation of these coefficients must also follow an approach where the coefficients of the interaction term are considered jointly with the coefficient of the stand-alone agent usage, but also on what their current and lagged motifs are. The overall conclusions to draw from the interaction term coefficients is that Slow Drawdowns have a higher average deposit amount when they use agents, above and beyond what can be gathered from the account expressing a Slow Drawdown motif and used an agent that period.

When we make pairwise comparisons β_5, Slow Drawdowns continue to be the motif with a statistically different interaction effect compared to any of the other three motifs (Table 27). There is no discernible difference between any of the other three motifs.

Prob($\beta_{5,X} = \beta_{5,Y}$)			
Y	X		
	Fast Drawdown	Accumulator	Slow Drawdown
Dump-and-Pull	0.178	0.323	0.000 ***
Fast Drawdown		0.783	0.100 *
Accumulator			0.033 **
Pairwise comparisons are being made across the coefficients of motifs X and Y *, ** and *** represent significance at 10%, 5% and 1% respectively.			

Table 27. Pairwise Comparison of Interaction Coefficients, Avg Deposit Amt

We thus find that deposit amounts are significantly correlated with previous amounts, and with agent usage during the current segment. Current and previous motifs show differentiated behavior amongst themselves, as well as additional distinction through motif-agent interactions.

The Hansen test tells us that our instruments are correlated with the error term. Let us explore this issue closely as it has the potential to impact the validity of our entire specification. The error terms seem to be fairly normally distributed. The histogram of errors with a normal density plot superimposed demonstrates a slight skew Figure 37 (a). The P-P plot is quite linear, suggesting normality in the mid-range of the errors (Figure 37 (b)). There is some disturbance in both the lower tail evident in the Q-Q plot (Figure 37 (c)), though nothing egregious. The residuals can therefore be considered to be approximately normal, with minor deviations.

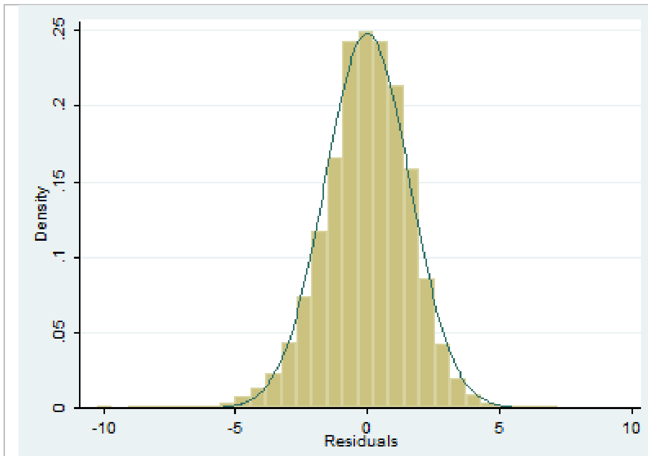

(a) Histogram for residuals, with scaled normal density plot

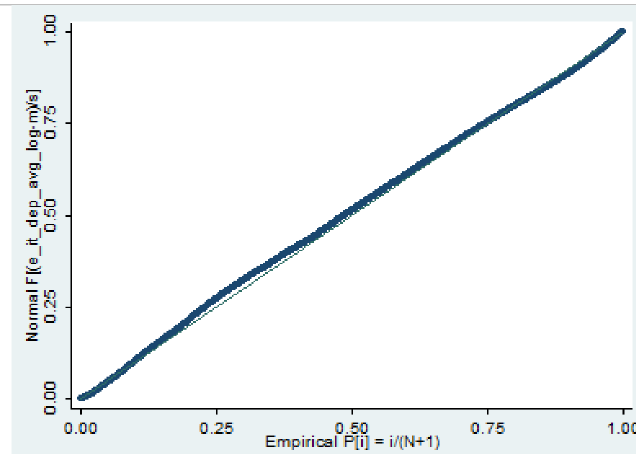

(b) Standardized normal probability plot (P-P)

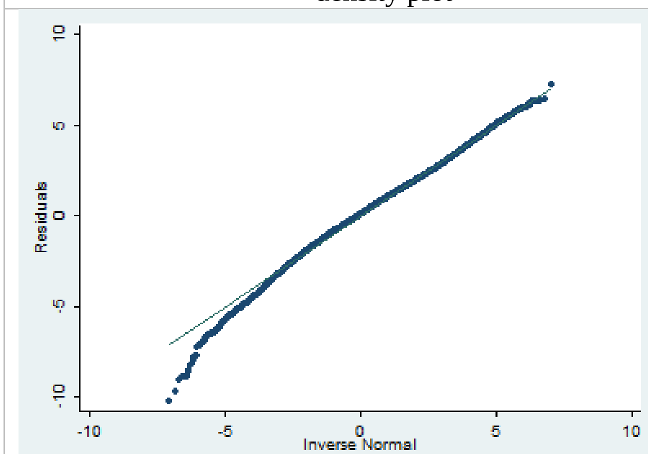

(c) Quintile plot of errors to normal distribution (Q-Q)

Figure 37. Error Plots, Average Deposit Amount

Since deposit amounts are considered to be endogenous and we expect correlation with the first lagged values, we test residuals against the second lagged value of average deposit amounts. The upward slope is quite evident in Figure 38. Regressing the second lagged logged average deposit amount to residuals gives the following, with an R^2 of 0.2178:

$$e_{it} = -5.228 + 0.592 * \log(avg\_deposits_{i,t-2})$$

Both the constant and coefficient are statistically significant at the 1% level, and the coefficient on *avg_deposits*$_{i,t-2}$ confirms the upward slope. Both the intercept and the slope are starker than what we have seen thus far. Adding additional lags does not instrument this correlation away. This implies that the second lagged deposit frequency is correlated with the error term.

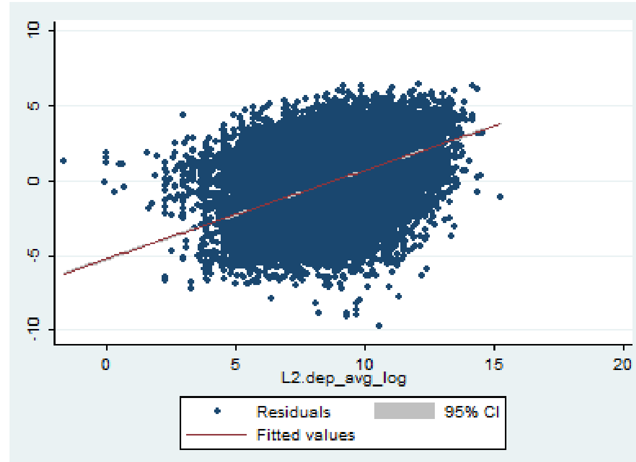

Figure 38. Residuals plotted against second lag of log(Avearge Deposit Amount)

We make the same case as for deposit that the effect of the instrumented variable in the form of the second lag of deposit frequencies will be overestimated, and must be mentally adjusted towards zero. To the extent that deposit amounts are serially correlated and the same bias transports over to the first lag, we can also mentally adjust the coefficients for the lagged deposit amounts towards zero. Our guesses for what this omitted variable could be are no different than what we discussed for balances, and are not rehashed.

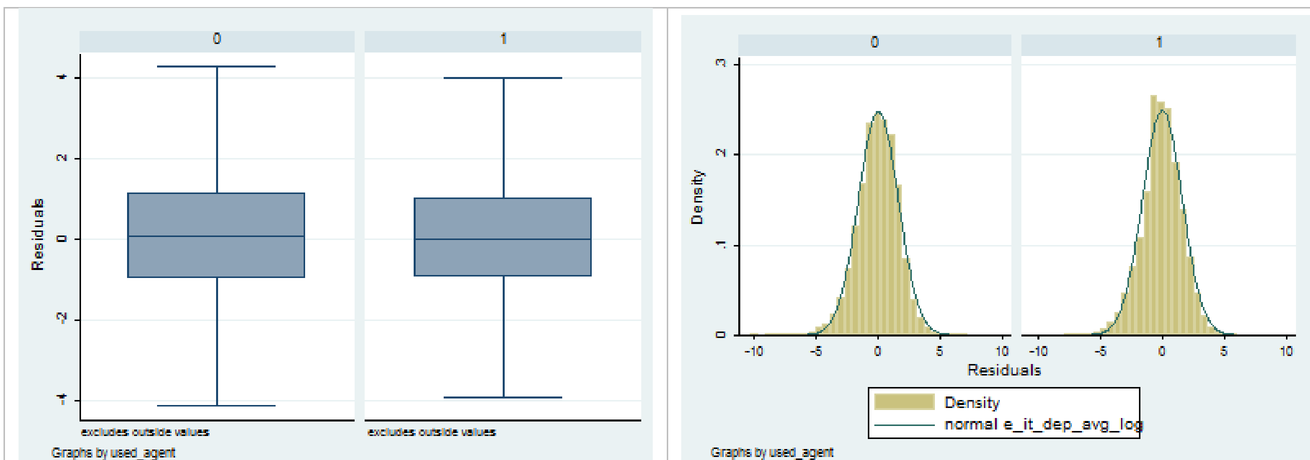

(a) Boxplot (b) Histogram with normal density plot

Figure 39. Residual Plots by Agent Usage

Since agent usage is considered to be exogenous and instrumented with itself alone, we can compare the residuals with the current agent usage. It seems that the errors are centered on zero from

the boxplot in Figure 39 (a), and have a normal distribution in Figure 39(b). There therefore does not seem be an issue with errors as far as agent usage is concerned.

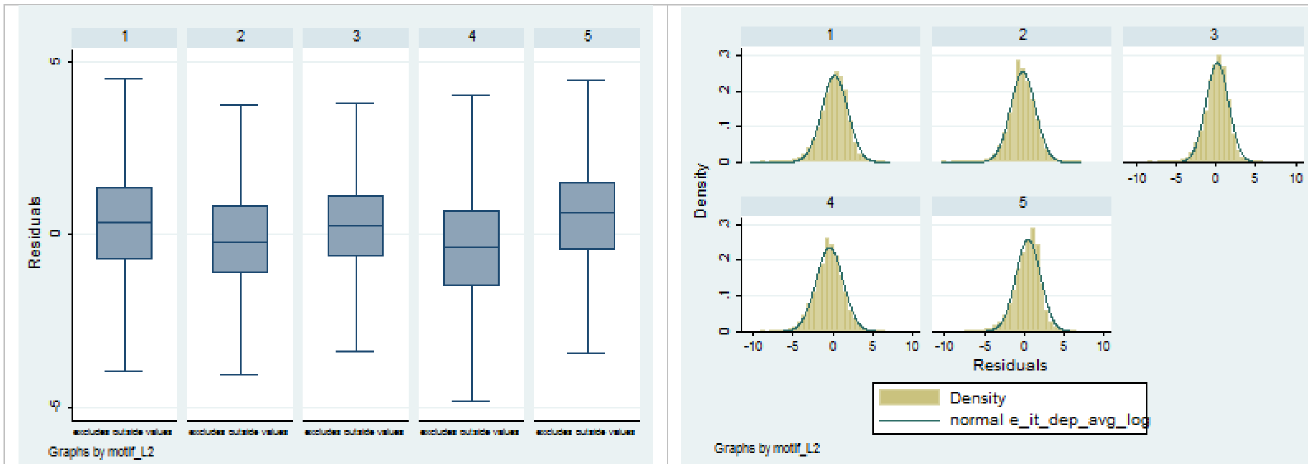

(a) Boxplot

(b) Histogram with normal density plot

Figure 40. Residual Plots against second Lagged Motif

Residuals are tested against motifs from two lags prior as motifs are suspected to be endogenous. The residuals look normally distributed Figure 40 (b), suggesting that we do not have a fundamental issue with our errors. There are however discernible deviations in Figure 40 (a), indicating the presence of possible omitted variable bias that is manifesting through correlation with this included covariate. Relying on the same logic presented earlier, we can conclude that the second lagged motif terms used as instruments are upward biased, and must be downward corrected, as must be the coefficients on current and first lagged motifs. In other words, the increases in average deposit amounts may not be as large as they seem for motifs that sport such an increase.

Number and Average Amount of Withdrawals

It continues to not be possible to remove persistent autocorrelation while exploring agents or motifs alone, as per specifications noted as (1) and (2) at the beginning of our discussion on average account balances. We will therefore be ignoring the two that were "agents only" and "motifs only" as we explore withdrawal dynamics too, and discuss the specifications for agents and motifs (3), and agents, motifs and their interaction (4) only.

Withdrawals	**Number**		**Average Amount**	
Explanatory Variables	(3)	(4)	(3)	(4)
Lagged log(withdrawal metric)	0.114 *** (0.029)	0.113 *** (0.030)	– 0.041 * (0.023)	– 0.033 (0.024)
Agent Used	– 0.099 *** (0.024)	0.353 (0.512)	– 0.265 *** (0.023)	– 1.025 ** (0.519)
M2: Dump-and-Pull	– 1.867 *** (0.261)	– 1.583 *** (0.338)	0.584 ** (0.282)	0.342 (0.368)
M3: Fast Drawdown	– 1.361 *** (0.352)	– 1.156 *** (0.425)	0.800 ** (0.358)	0.039 (0.456)
M4: Accumulator	– 2.961 *** (0.430)	– 3.080 *** (0.519)	1.797 *** (0.369)	1.124 ** (0.490)
M5: Slow Drawdown	– 0.478 (0.340)	– 0.198 (0.429)	0.591 (0.364)	0.053 (0.446)
Lagged M2: Dump-and-Pull	– 0.159 (0.235)	– 0.134 (0.238)	0.088 (0.222)	0.167 (0.229)
Lagged M3: Fast Drawdown	0.005 (0.343)	0.059 (0.347)	– 0.256 (0.274)	– 0.182 (0.281)
Lagged M4: Accumulator	0.257 (0.444)	0.286 (0.452)	1.611 *** (0.417)	1.620 *** (0.432)
Lagged M5: Slow Drawdown	0.429 (0.345)	0.479 (0.352)	– 0.287 (0.345)	– 0.263 (0.351)
Agent Used × M2		– 0.656 (0.656)		0.253 (0.675)
Agent Used × M3		– 0.617 (0.849)		2.350 *** (0.881)
Agent Used × M4		0.393 (0.930)		0.979 (0.993)
Agent Used × M5		– 0.650 (0.715)		1.498 ** (0.706)
Intercept	3.244 *** (0.316)	3.007 *** (0.380)	7.834 *** (0.338)	8.080 *** (0.376)
AR(1)	0.000	0.000	0.000	0.000
AR(2)	0.774	0.808	0.075	0.206
Hansen test	0.003	0.003	0.000	0.001

Wald χ^2	0.0000	0.0000	0.0000	0.0000
F-test for Motif Joint Sig.	0.0000	0.0000	0.0000	0.0276
F-test for Joint Sig., Lagged Motifs	0.0015	0.0013	0.0000	0.0000
F-test for Joint Sig., Interaction	-	0.5571	-	0.0090
No. of observations	217,826	217,826	217,826	217,826
No. of accounts (groups)	22,606	22,606	22,606	22,606
No. of instruments	196	196	196	196
Standard errors are in parentheses. *, ** and *** represent significance at 10%, 5% and 1% respectively. AR(1) – 1st order autocorrelation test; AR(2) – 2st order autocorrelation test (1) and (2) use all available lags as instruments; (3) and (4) use second lags only.				

Table 28. System GMM Estimators for Frequency and Amount of Withdrawals

First, we look at the effect of motifs and agent usage on the number of withdrawals. Figure 41 provides a more visual representation of the coefficients for (4), making it easier to compare magnitude of impact.

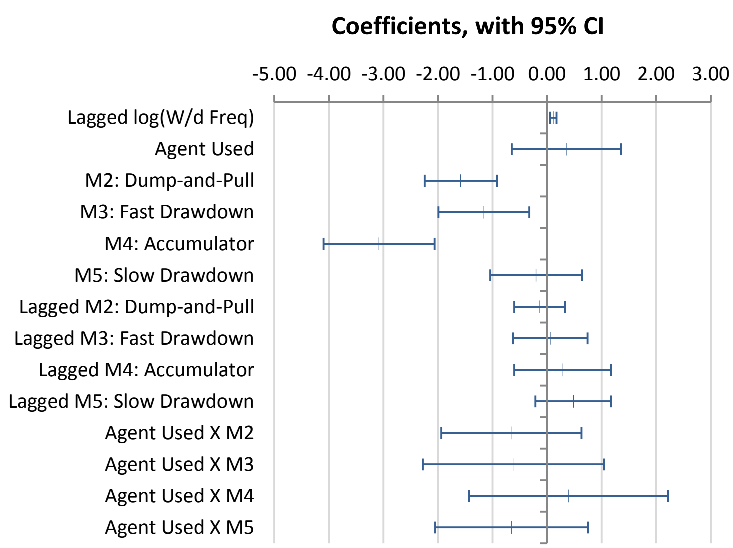

Figure 41. Coefficients with 95% CI for Withdrawal Frequency

Second-order aurocorrelation for number of withdrawals is mitigated through GMM estimators in both (3) and (4). The Hansen test is weak, suggesting that our instruments continue to be correlated with the error term. The number of instruments, 196, is much smaller than the number of accounts, at 22,606, implying that we do not have a "too many instruments" problem. Let us look at each component of our model.

- log(*number of withdrawals*$_{i,t-1}$) : An approximately 10% increase in the number of withdrawals can be expected from the previous period's figures. This implies that the number of withdrawals increases over time.
- *used_agent*$_{i,t}$: Specification (3) suggests that the overall increase in frequency is about 11% when an agent is used. The coefficient of (4) is not statistically different than zero for Sustained Balances; the other motifs cannot generally be looked at independently of the interaction term. Discussion is deferred to the appropriate section below.
- *motif*$_{i,t}$: Three of the four motif coefficients are statistically significant, and the F-test for joint significance, which has the null hypothesis that all coefficients on motifs are statistically indistinguishable from zero, is significant at the 1% level. This suggests that at least some of the motifs have statistically significantly different withdrawal frequencies. We find Dump-and-Pull, Fast Drawdown and Accumulator accounts have lower withdrawal frequencies than Sustained Balances, while that of Slow Drawdowns is indistinguishable from it.
- *motif*$_{i,t-1}$: None of the four lagged motif coefficients are statistically significant for lagged motif. Interestingly, the F-test for joint significance for these lagged motifs suggests statistical significance, with a p-value of 0.0013. We suspect this is simply a result of the large sample size and unbalanced dataset which causes discernible differences in variances even though the coefficients are not significant.
- *motif*$_{i,t}$**used_agent* : None of the four interaction terms are statistically significantly different from that of Sustained Balances, and this time this result is confirmed by the F-test for joint significance. The overall conclusion to draw from the interaction term

coefficients is that there is no motif-agent use combination that stands out as being different from any of the other ones.

There is no statistically significant pairwise difference in the interaction effect when any of the other motifs are concerned either, as can be seen in Table 29 below.

$Prob(\beta_{5,X} = \beta_{5,Y})$			
Y	X		
	Fast Drawdown	Accumulator	Slow Drawdown
Dump-and-Pull	0.957	0.185	0.993
Fast Drawdown		0.315	0.967
Accumulator			0.180
Pairwise comparisons are being made across the coefficients of motifs X and Y *, ** and *** represent significance at 10%, 5% and 1% respectively.			

Table 29. Pairwise Comparison of Interaction Coefficients, Withdrawal Frequency

We thus find that withdrawal frequencies are significantly correlated with previous frequencies, but not with agent usage during the current segment. Current motifs show differentiated behaviors amongst themselves, but previous motifs do not. Additional distinction is not available through motif-agent interactions.

The Hansen test tells us that our instruments are correlated with the error term. Let us explore this issue closely as it has the potential to impact the validity of our entire specification. The error terms seem to be fairly normally distributed. The histogram of errors with a normal density plot superimposed demonstrates a disturbance, but it does not impede the overall shape akin to a normal distribution Figure 42 (a). The P-P plot is quite linear, suggesting normality in the mid-range of the errors (Figure 42 (b)). There is some disturbance in both the upper and lower tails evident in the Q-Q plot (Figure 42 (c)), though it seems to be fairly minimal. The residuals can therefore be considered to be approximately normal, with minor deviations.

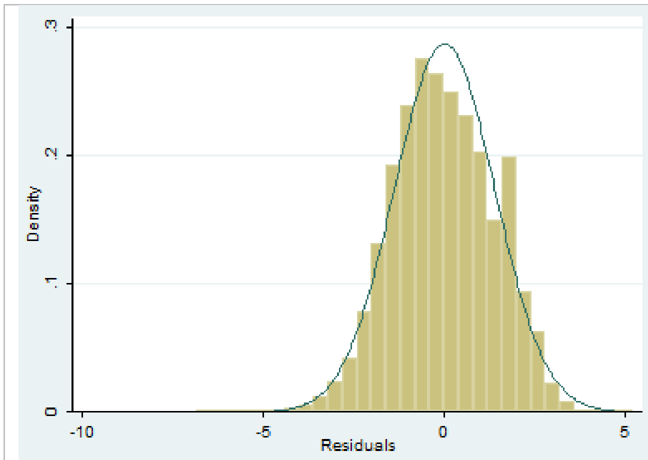

(a) Histogram for residuals, with scaled normal density plot

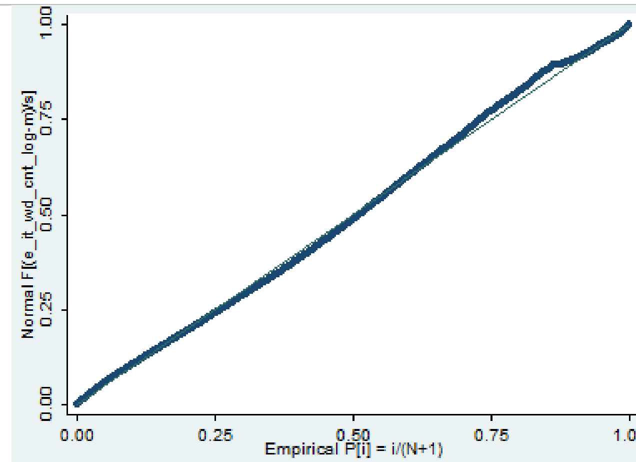

(b) Standardized normal probability plot (P-P)

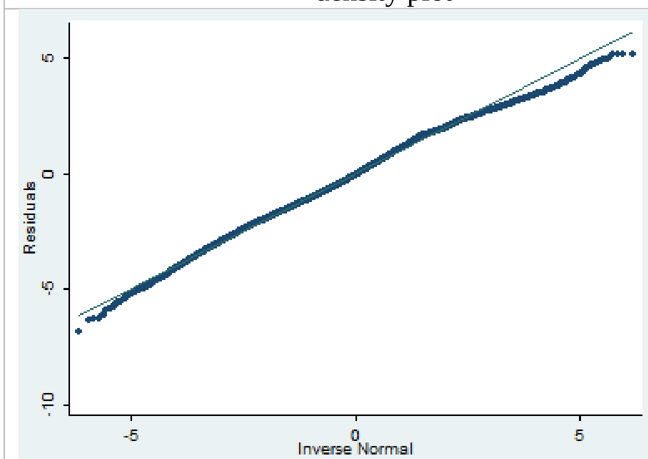

(c) Quintile plot of errors to normal distribution (Q-Q)

Figure 42. Error Plots, Withdrawal Frequency

Since number of withdrawals is considered to be endogenous and we expect correlation with the first lagged values, we test residuals against the second lagged value of withdrawal frequencies. The upward slope is quite evident in Figure 43. Regressing the second lagged logged withdrawal frequency to residuals

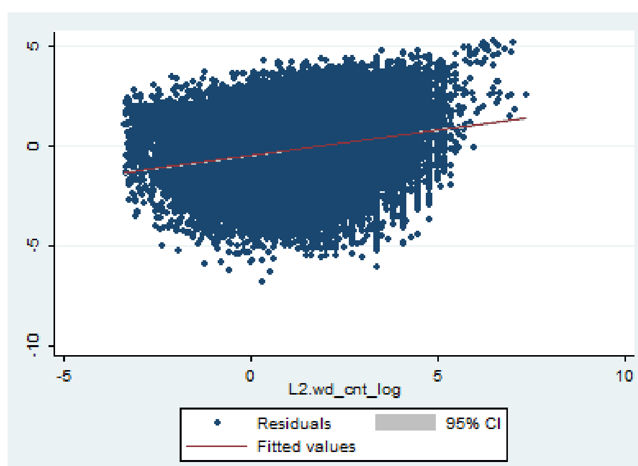

Figure 43. Residuals plotted against second lag of log(Number of Withdrawals)

gives the following, with a very weak R^2 of 0.0464:

$$e_{it} = -0.491 + 0.254 * \log(num\_withdrawals_{i,t-2})$$

Both the constant and coefficient are statistically significant at the 1% level, and the coefficient on $num\_withdrawals_{i,t-2}$ confirms the upward slope. Adding additional lags does not instrument this correlation away. This tells us that the second lagged withdrawal frequency is correlated with the error term.

We suspect omitted variable bias here too, and can make the case that the effect of the instrumented variable in the form of the second lag of withdrawal frequencies will be overestimated, and must be mentally adjusted towards zero. To the extent that withdrawal frequencies are serially correlated and the same bias transports over to the first lag, we can also mentally adjust the coefficients for the lagged withdrawal frequency towards zero. Our guesses for what this omitted variable could be are no different than what we discussed for balances and deposits, and are therefore not rehashed.

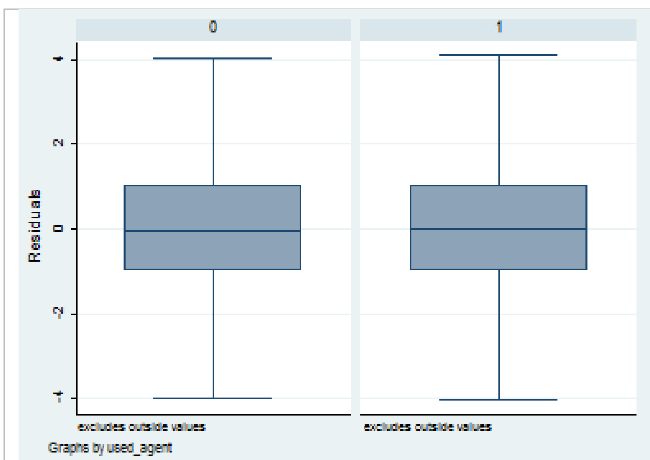

(a) Boxplot

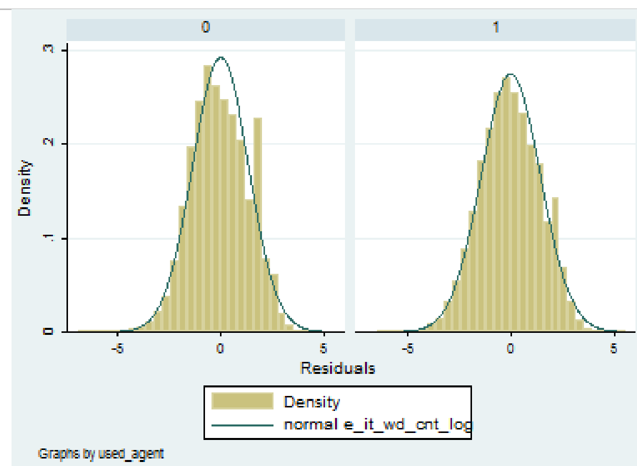

(b) Histogram with normal density plot

Figure 44. Residual Plots by Agent Usage

Since agent usage is considered to be exogenous and instrumented with itself alone, we can compare the residuals with the current agent usage. It seems that the errors are centered on zero from

the boxplot in Figure 44(a), and have a normal distribution despite the tell-tale spike in Figure 44Figure 34 (b). There therefore does not seem be an issue with errors as far as agent usage is concerned.

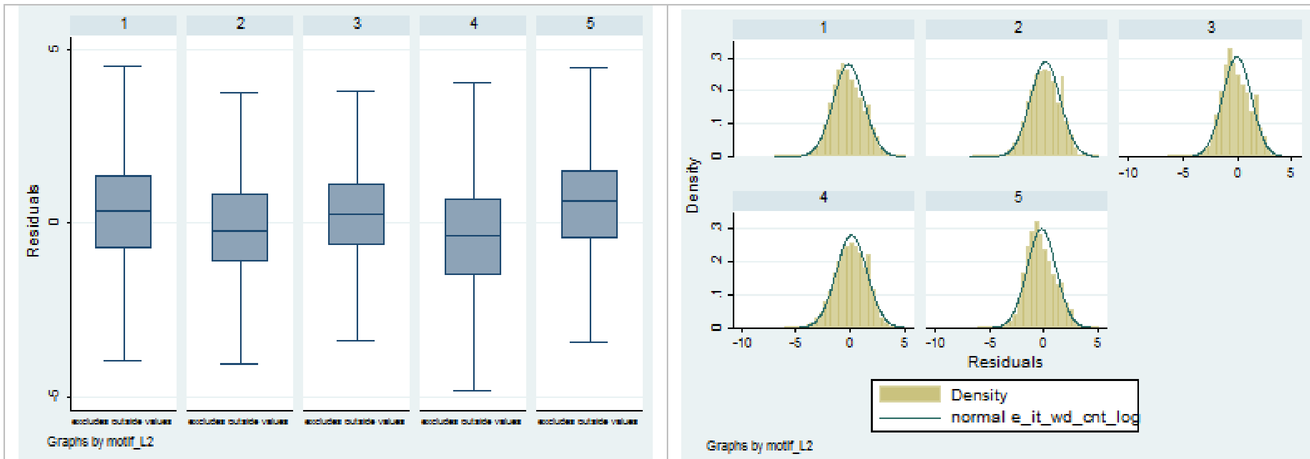

(a) Boxplot

(b) Histogram with normal density plot

Figure 45. Residual Plots against second Lagged Motif

Motifs are suspected to be endogenous, and therefore residuals are tested against motifs from two lags prior. We can see deviations in the boxplots for residuals for all second lagged motifs Figure 45 (a). The residuals look normally distributed Figure 45 (b). Therefore there seems to be some issues with residual distribution as far as motifs are concerned.

We can extend the argument we presented earlier for the lagged outcome variable to motifs too, in that the correlation of both the outcome variable and the included independent variable must be in the same direction with the omitted independent variable as they are but lagged versions of each other, with the current values becoming the lagged values in due course. Since the correlation must be either positive for both, or negative for both, the estimate for the lagged motif instruments are overestimated, and must be corrected downward, towards zero. In so far as lagged motifs of the current run are current motifs in some other run, we can extrapolate these results to say that we must also downward correct the coefficients of *motif*$_t$, and *motif*$_{t-1}$. In other words, the decreases in withdrawal frequencies are not as large as they seem for motifs that sport such a decrease.

Next, we look at the effect of motifs and agent usage on the average size of withdrawals. Figure 46 provides a more visual representation of the coefficients for (4), making it easier to compare magnitude of impact.

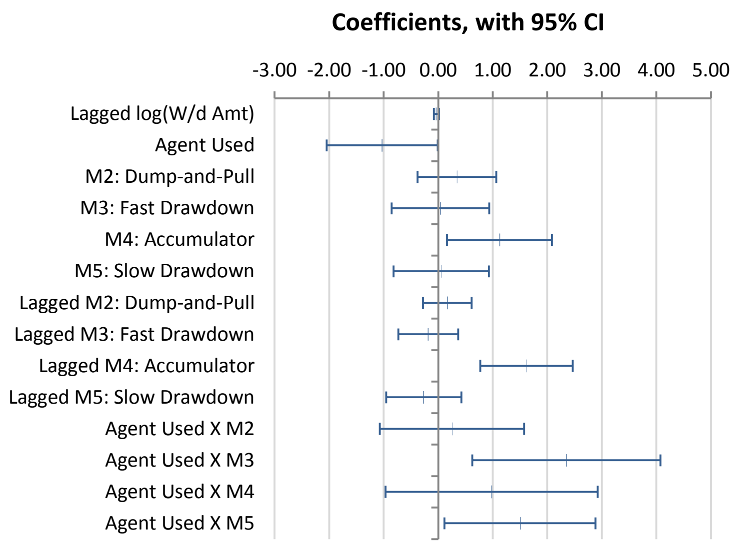

Figure 46. Coefficients with 95% CI for Avg Withrawal Amt

Second-order aurocorrelation for average withdrawal amounts is mitigated through GMM estimators in both (3) and (4). The Hansen test is weak, suggesting that our instruments continue to be correlated with the error term. The number of instruments and accounts are the same as for deposit frequencies, implying that we do not have a "too many instruments" problem. Let us look at each component of our model.

- log(*avg_withdrawals*$_{i,t-1}$) : When no agent interaction is considered, (i.e. specification (3)), there is a 4% decrease in the average withdrawal amount between periods, though this is only significant at the 10% level. When the interaction term is brought in, the significance of the coefficient disappears. We conclude that there is no effective change in the average withdrawal amounts, over time.

- *used_agent*$_{i,t}$: Both specifications, with or without interaction, suggest a reduction of average withdrawal amount is associated with agent usage. Specification (4) refers to Sustained Balances only though; we note the effect on other motifs below.
- *motif*$_{i,t}$: Only one of the four motif coefficients are statistically significant, and the F-test for joint significance, which has the null hypothesis that all coefficients on motifs are statistically indistinguishable from zero, is significant at the 1% level. This suggests that Accumulators have a higher average withdrawal amount compared to the other four, which in themselves are indistinguishable from the Sustained Balance accounts.
- *motif*$_{i,t-1}$: Accumulators also show a statistically significant higher average withdrawal amount compared to the other four motifs, for lagged motif coefficients. The F-test for joint significance is significant, at the 1% level. This suggests that Accumulators consistently have higher average withdrawal amounts.
- *motif*$_{i,t}$**used_agent* : Fast and Slow Drawdown accounts show increased average withdrawal amounts when combined with agent usage. When combined with the coefficients on *used_agents*$_{it}$ alone, it implies that the net change in average withdrawal amount is positive for Fast and Slow Drawdowns, and negative for the other three motifs, when compared to cases where no agent is used.

When pairwise comparisons with other motifs are considered, Fast and Slow Drawdowns show a statistically different interaction effect compared to Dump-and-Pulls (Table 30). There are no other differences between other motifs.

Prob($\beta_{5,X} = \beta_{5,Y}$)			
Y	X		
	Fast Drawdown	Accumulator	Slow Drawdown
Dump-and-Pull	0.008 ***	0.377	0.027 **
Fast Drawdown		0.224	0.328
Accumulator			0.552
Pairwise comparisons are being made across the coefficients of motifs X and Y *, ** and *** represent significance at 10%, 5% and 1% respectively.			

Table 30. Pairwise Comparison of Interaction Coefficients, Avg Withdrawal Amt

We thus find that current withdrawal amounts are significantly correlated with previous frequencies, and with agent usage during the current segment. Current motifs show differentiated behaviors amongst themselves, as do previous motifs. Additional distinction is available through motif-agent interactions.

The Hansen test tells us that our instruments are correlated with the error term. Let us explore this issue closely as it has the potential to impact the validity of our entire specification. The error terms seem to be fairly normally distributed. The histogram of errors with a normal density plot superimposed demonstrates a slight skew (Figure 47(a)). The P-P plot is quite linear, suggesting normality in the mid-range of the errors (Figure 47 (b)). There is noticeable disturbance in both the lower tail evident in the Q-Q plot (Figure 47 (c)). The residuals can therefore be considered to be approximately normal for the most part, with discernible deviations in its lower tails.

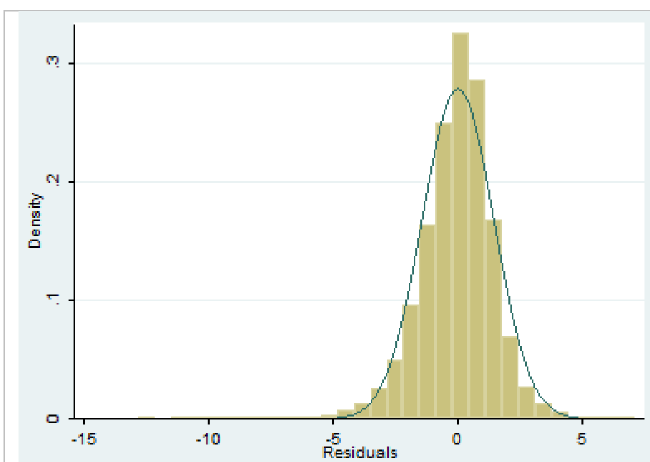

(a) Histogram for residuals, with scaled normal

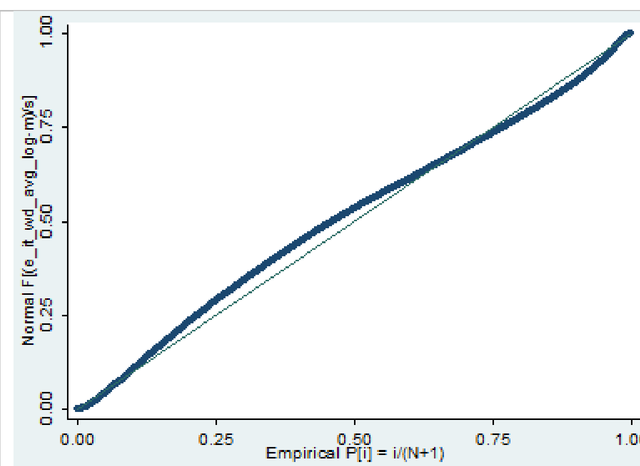

(b) Standardized normal probability plot (P-P) density plot

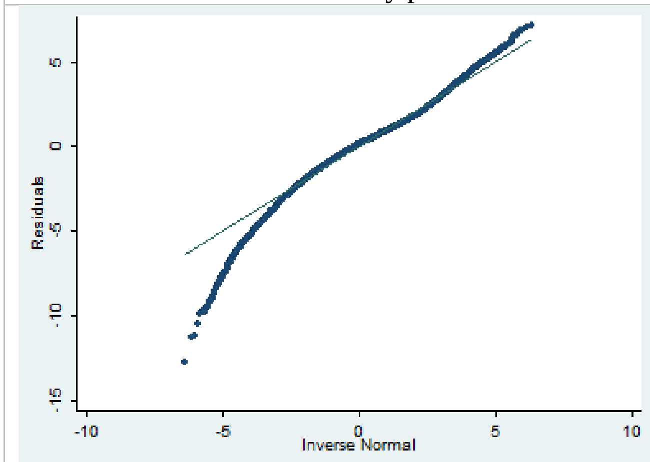

(c) Quintile plot of errors to normal distribution (Q-Q)

Figure 47. Error Plots, Withdrawal Amount

Since number of deposits is considered to be endogenous and we expect correlation with the first lagged values, we test residuals against the second lagged value of average withdrawal amounts. The upward slope is quite evident in Figure 48. Regressing the second lagged logged balance to residuals gives following, with a low R^2 of 0.1767:

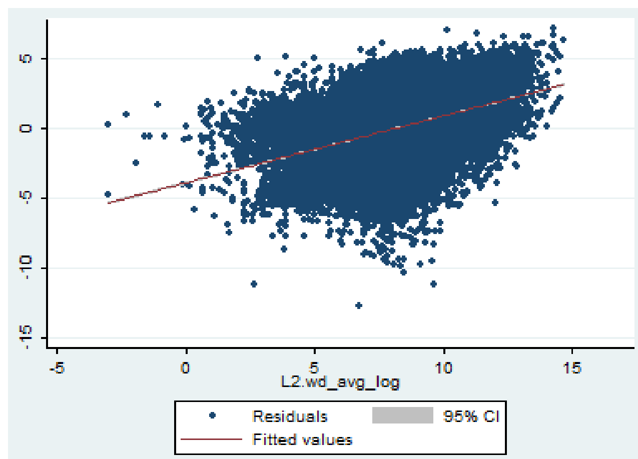

Figure 48. Residuals plotted against second lag of log(Average Withdrawal Amount)

$$e_{it} = -3.917 + 0.480 * \log(avg\_withdrawals_{i,t-2})$$

The coefficient is statistically significant at the 1% level, and the coefficient on $avg\_withdrawalss_{i,t-2}$ confirms the upward slope. Both the intercept and the slope are starker than what we have seen thus far. Adding additional lags does not instrument this correlation away. This leaves us with the possibility that the second lagged deposit frequency is correlated with the error term, though not any of the other covariates.

Here too, we must make the case that the effect of the instrumented variable in the form of the second lag of withdrawal amounts will be overestimated, and must be mentally adjusted towards zero. To the extent that average withdrawal amounts are serially correlated and the same bias transports over to the first lag, we can also mentally adjust the coefficients for the lagged withdrawal amounts towards zero. Our guesses for what this omitted variable could be are no different than what we discussed for balances, and are therefore not rehashed.

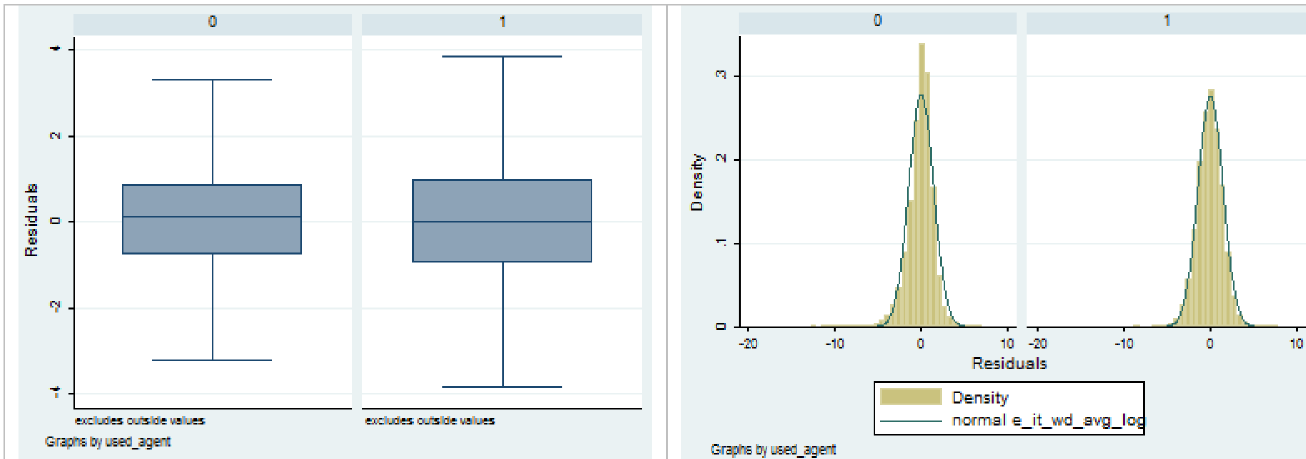

(a) Boxplot (b) Histogram with normal density plot

Figure 49. Residual Plots by Agent Usage

Since agent usage is considered to be exogenous and instrumented with itself alone, we can compare the residuals with the current agent usage. It seems that the errors are centered on zero from the boxplot in Figure 49 (a), and have a normal distribution in Figure 49 (b). There therefore does not seem be an issue with errors as far as agent usage is concerned.

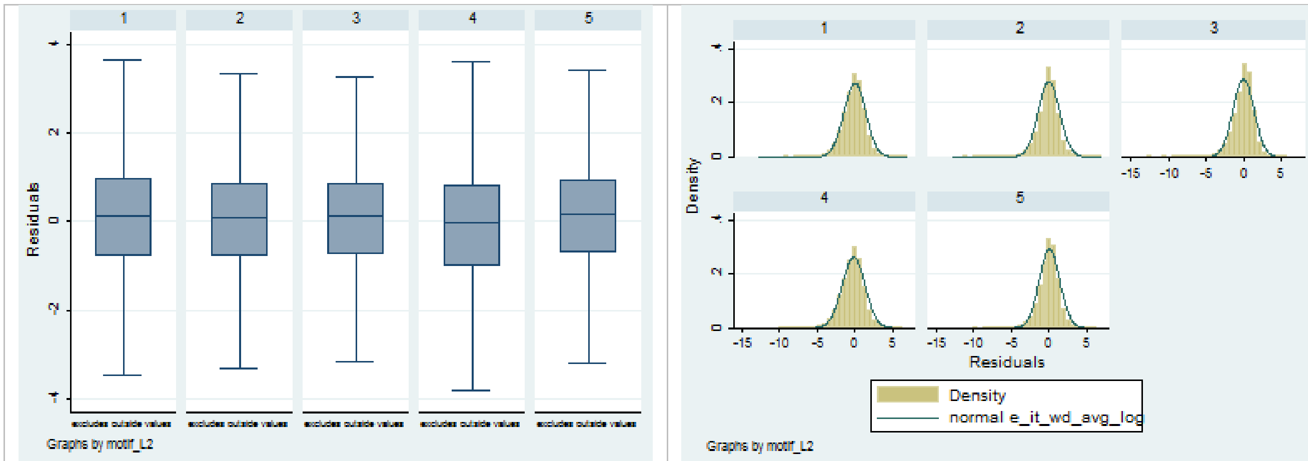

(a) Boxplot (b) Histogram with normal density plot

Figure 50. Residual Plots against second Lagged Motif

Residuals are tested against motifs from two lags prior as motifs are suspected to be endogenous. The residuals look normally distributed Figure 50 (b), suggesting that we do not have a fundamental issue with our errors. There are also essentially centered around zero, as seen in the boxplots in Figure 50 (a), indicating very little disturbances in errors. Given the minimal nature of disturbances, we infer that there is little by way of omitted bias issues we need to be concerned with, as far as motifs are concerned.

Summary of Findings

We condense the findings from using Arellano-Bond system GMM Estimators on balances, deposits and withdrawals in Table 31 below. In it, we summarize answers to the following four questions asked with respect to the specification we use:

- Does the outcome variable increase or decrease over time?
- Are outcome variable levels statistically significantly different depending on whether an agent has been used?

- Are outcome variable levels statistically significantly different depending on what motif a segment is associated with in the present and immediate past?
- Are outcome variable levels statistically significantly different when agent usage is taken into consideration along with motifs?

In each of these cases, we note the positive bias where it exists by an upward arrow (↑) as superscript, where coefficients have been overestimated due to correlation with the omitted independent variables. We did not identify any downward bias for any of the dependent or independent variables.

It is worth nothing that such positive bias implies that we must mentally downward adjust our estimators, potentially to the point where it is no longer significant at all. The "Motif" column indicates how many of the remaining four motifs show a statistically significant difference compared to Sustained Balances.

Outcome Variable	Change Over Time	Motif	Agent Usage + Motif
Average Balance	Decreases$^{\uparrow}$	3 of 4	No Change
Number of Deposits	Increases$^{\uparrow}$	2 of 4$^{\uparrow}$	Increase: Sustained Balance, Accumulators Decrease: Dump-and-Pull, Fast and Slow Drawdown
Avg. Deposit Amount	Increases$^{\uparrow}$	2 of 4$^{\uparrow}$	Increase: Slow Drawdown No Change: Other Four
Number of Withdrawals	Increases$^{\uparrow}$	3 of 4$^{\uparrow}$	No Change: All Five
Avg. Withdrawal Amount	No Change$^{\uparrow}$	1 of 4$^{\uparrow}$	Increase: Fast and Slow Drawdown Decrease: Sustained Balance, Dump-and-Pull and Accumulator

Table 31. Summary Interpretation of AB-GMM Estimators ($^{\uparrow}$ indicates positive bias)

First, let us glance at findings that flow naturally from our specification, even though they are not directly related to our interest in agent usage. Lagged outcome variable estimators, albeit overestimated, suggest a secular increase in number of deposits, average amount of deposits, and number of withdrawals over time, and a secular decrease of balances over time, with no change in the average amount of withdrawals. The changes in outcome variables when segregated by motif, where they exist, are expected and fairly consistent with statistics presented earlier during the motif discovery process.

The agent impact is quite interesting. The information presented in Table 31 is re-presented below in Table 32 to highlight this focusing only on the interaction of motifs with agents. A positive interaction effect that is significant at at least the 5% level is noted with a green plus – this implies that the outcome variable has a statistically significant positive correlation with agent usage and motif, above and beyond the effect from having used an agent or displaying a certain motif alone. Similarly, a negative statistically significant interaction effect, significant at the 5% level, is noted with a red minus. We note 5% as the cut-off p-value because coefficients of interaction effects are either significant at the 1% or 5% level, or not significant at all – there are none that are significant at the 10% level (Table 22, Table 25, and Table 28).

Outcome Variable	**M1: Sustained Balance**	**M2: Dump-and-Pull**	**M3: Fast Drawdown**	**M4: Accumulator**	**M5: Slow Drawdown**
Average Balance	•	•	•	•	•
Number of Deposits	+	−	−	+	−
Avg. Deposit Amount	•	•	•	•	+
Number of Withdrawals	•	•	•	•	•
Avg. W/d Amount	−	−	+	−	+

Table 32. Direction of Agent-Motif Interaction Effect, Stat. Sig. Coefficients Only

This suggests that there does not seem to be any difference in average balances, and the number of withdrawals associated with agent usage, regardless of the type of user. Accumulators and Sustained Balances increase the number of deposits, and decrease the average withdrawal size when given access to an agent. Dump-and-pulls show a decrease in number of deposits and in the average amount of withdrawals. It is difficult to discuss what is happening given that for most of the motifs, we do not know the direction of movement for three of the five outcome variables.

Indeed, it is difficult to ascertain if these results make sense at all. For example, if the average balance for Sustained Balances has not changed, despite a higher number of deposits and smaller withdrawals, then the average deposit amounts must have decreased, or the number of withdrawals increased, or both, even though this is not picked up by our AB-GMM runs. Similarly, if the average balance for Fast Drawdowns has not changed, then having fewer deposits but higher withdrawal amounts associated with agent usage makes sense only if the deposit size increases or the number of withdrawals decreases. Changes to deposit amounts may have a different implication than changes to withdrawal frequencies, depending on the direction of change, making it difficult to ascertain what is happening.

To help us navigate these uncertainties, we re-present Table 32 below in Table 33 and include all the point estimates of the coefficients that are not statistically significant at the 5% level and were initially left out. We consider these point estimates for average balances, withdrawal frequencies and deposit amounts as being indicative of their general tendency of how they may change with agent usage, and find that this helps us understand what is happening within the motifs even though they lack the statistical certainty of deposit frequencies and withdrawal amounts.

Positive point estimates are leaner grey pluses, and negative point estimates are leaner grey minuses. Coefficients which are virtually indistinguishable from zero are represented by a grey circle.

Outcome Variable	M1: Sustained Balance	M2: Dump-and-Pull	M3: Fast Drawdown	M4: Accumulator	M5: Slow Drawdown
Average Balance	−	○	+	○	−
Number of Deposits	**+**	**−**	**−**	**+**	**−**
Avg. Deposit Amount	−	−	○	−	**+**
Number of Withdrawals	+	−	−	+	−
Avg. Withdrawal Amount	**−**	**−**	**+**	**−**	**+**

Table 33. Direction of Agent-Motif Interaction Effect, All Coefficients

Table 33 provides us with a much clearer picture of what is happening. Both Sustained Balances and Accumulator increase deposit and withdrawal frequencies, and reduce deposit and withdrawal amounts. This implies that they interact with their account more often when using agents, but do so in smaller amounts. The unchanging balances for Accumulators suggest that agent usage is not associated with the ability to save a different lump-sum amount compared to when agents are not used. The reduction in balances for Sustained Balances, albeit not statistically significant, could be a result of either earlier withdrawals, a smaller amount funds handed during that segment, or both. These motifs therefore represent more vigorous engagement with the savings account in the presence of an agent, with a possible marginal neutral or negative impact on balances.

We see the opposite effect for Fast and Slow Drawdowns, where the frequency of deposit and withdrawals go down, but the amount of withdrawals go up for both motifs, while the amount of deposits goes up for Slow Drawdowns. This result is internally consistent for both, in that smaller deposit and withdrawal frequencies can combine with larger amounts of both to result in negligible changes in balance. Given that we pictured Slow Drawdowns as the motif paying off regular expenses with initial top-up(s), it suggests that these interactions have become more lumpy for this motif, though we cannot offer a good explanation. Fast Drawdowns are characterized by a large initial withdrawal followed by a residual balance, so fewer, larger withdrawals with a tendency for a higher balance might mean that more of the initial top-up is taken out in the first withdrawal, but that it is done later than when agents are not used.

Dump-and-Pulls seem most adversely affected motifs – frequencies and amounts of both deposits and withdrawals are lower when agents are used. This may simply suggest that agent usage is not conducive for these very short term back-to-back deposit and withdrawal cycles.

We will fully explore the implications of these findings in the next Chapter, within the greater context of expected and observed savings behaviors.

We end this section with a note on causality. Given the distinct lack of other explanatory variables, we will not make the case that what we offer is a causal model behind agent usage and the various outcome variables, in that we will not suggest that number of deposits went up or average withdrawal amounts went down *because* of agent usage. This is primarily because we suspect the presence of time-variant omitted variable bias, despite having the following mitigating factors:

- All fixed effects are accounted by through differencing. Thus, time-invariant omitted variables cannot confound our estimators.
- Time-variant omitted variables have been shown to cause overestimation of estimators for lagged outcome variables in all cases and some motifs in some cases, but not agent usage.
- We are fortunate to have a rather large dataset – about 20,000 accounts offering between 176,000 and 228,000 segments that are regressed over using 196 instruments, offering significant statistical power. While this does not have any direct implication on causality, it does make it highly unlikely for spurious correlations to occur.

Conclusion

In our literature review, we explored how the poor can and do save, but the options for saving in formal financial instruments are often few or unattractive. This thesis took a detailed analytical look at the no-frills Ordinary Savings accounts offered by Bank A in Kenya to explore existing savings behavior. In doing so, it offered quantitative support for anecdotal frameworks that are prevalent in the savings literature in the financial inclusion space, formalized five specific behavioral patterns, and identified the changes of behavior associated with agent usage. In this concluding section, we summarize these findings and reflect on the implications of said findings mostly on individuals and FIs, but also with some thought spared for practitioners and policy makers.

Key Findings

We identified five major recurring patterns, or motifs, that manifest themselves in the voluntary savings accounts, of which four can be characterized as representing saving. **Accumulators** are who we generally picture as savers, putting away small amounts of money over time to build a lump sum, before drawing it down. If additional withdrawals do occur before the final drawdown, they do not disrupt the generally increasing nature of balances. The average accretive saving account takes about a month to complete a cycle from starting to build savings to withdrawing completely. This corresponds to what has been described as "saving up" by Rutherford, and tagged as "Type B" savings by GAFIS.

This segment is quite interesting because saving up in this manner is difficult in general, given the discipline required to persist in this behavior over time. It is doubly difficult for low-income individuals, who have to make choices between saving, investment and consumption, with pressing needs for funds cropping up often. That a significant share of individuals use these accounts to save deliberately over time signals success for these accounts in promoting that specific, successful savings behavior where individuals are able to enforce sufficient discipline on themselves to build up savings.

Three motifs fall under what is recognized as “saving down” behavior a la Rutherford. All of them initiate the savings cycle by depositing a lump-sum of money at the onset, which is then drawn down about three weeks after initiation, on average. **Sustained Balance** motifs hold on to most of the funds from the initial deposit, with some fluctuation as small deposits and withdrawals are conducted, and then begin withdrawing around the second half of the segment. This seems to represent situations where there is a mismatch of funds inflow and outflow, with clients using the account as the intermediate repository.

Slow Drawdown motifs seem to represent those who have a more deliberate, regular use of the account in mind to pay for periodic expenses with an initial deposit. The regularity of withdrawals is notable as it potentially signals a comfort with the use of the medium, trust in the asset, and availability on demand. Amongst the five motifs, Slow Drawdowns and Sustained Balances have the highest and second highest average deposit amounts respectively. Possible sources of initial funds could be salaries, a large sale at one’s place of business, or remittances received, though we cannot ascertain that information from available data. These two motifs probably do not represent clients who find formal savings unreliable, untrustworthy or inconvenient, as per the section *Why Don’t The Poor Save?*

Fast Drawdown motifs are somewhat different, conducting a sizeable withdrawal soon after the initial deposit and leaving residual funds in the account for the remainder of the period. We do not know if the remaining funds are a deliberate attempt to save away a small amount after a large expense, or simply the leftover for which an immediate use is not recognized. Deliberation would represent a tenacious desire to squirrel away funds when available, which would be proof of an individual's desire to save. The possibility of savings as an afterthought is perhaps less glamorous, but nevertheless has the same effect of residual savings.

The fifth motif, **Dump-and-Pull**, does not qualify to be called savings, as everything that is deposited is almost immediately withdrawn within four days for the average account. While such usage of the account as a mere pass-through could be seen as disappointing as far as savings promotion is concerned, especially considering the fact that two out of every five clients are using it thus, it is also a sign of user ingenuity and adaptability. An account intended for one purpose but with rules that also incentivize other uses has been coopted to facilitate those other uses. Since no-frills accounts have no deposit charges and withdrawal charges in limited situations and at much lower levels than other options, it is reasonable for individuals to use it as a pass-through. While this undoubtedly has utility for the client, this motif is quite expensive to support for the bank, as a result of the lack of float. We will revisit this issue of the business case for each account in more detail shortly.

In summary, Accumulators confirm the existence of accretive saving using a formal financial instrument, akin to what is achieved through discipline via piggy-banks and savings groups. Three motifs offer clarity on specific modes of "saving down" – saving the residual (Fast Drawdown), orderly depletion to pay regular expenses (Slow Drawdown) and as a repository of excess funds that is tapped into as needed (Sustained Balance). Lastly, the most voluminous completely transactional non-savers lurking within these savings accounts showed adaptation by clients to serve a financial need that was not originally intended.

Let us now consider the impact of using agents and how that affects motif behaviors. Note that we are discussing statistical significant correlations with agent usage that we interpret to signal the strong possibility of causation. In our interpretations, we rely primarily on coefficients which are statistically significant at the 5% level, but also take help from non-zero point estimates for coefficient that are not statistically significant to guide out interpretation.

Accumulators and Sustained Balance motifs *increased* the frequency of deposits, and *decreased* the average withdrawal amount when agents were used. There was also a tendency to increase the number of withdrawals, and decrease the average deposit amount. The increase in transaction frequencies is probably because of agent proximity that makes it more convenient for individuals to save more often, and access those funds more often too. A reduced transaction amount is consistent with balance level that change little or none at all, as it implies that the same amount of funds is being intermediated, but with greater granularity. Thus, it seems that *agent usage is associated with increasing engagement of individuals with these accounts, and doubling down on what they were doing all along* – maintaining an account balance through multiple deposits and withdrawals.

Agents seem to reinforce some of the mechanics of Slow and Fast Drawdown motifs, though they may not be quite so saving-oriented. Slow Drawdowns have fewer deposits of larger amounts, and withdraw in larger amounts, when using agents. These users seem to be doubling down on their use of the account as a repository to pay regular expenses from – *they top it up with more funds than they used to*, and also take out more than they used to.[3] *Agents therefore seem to discourage intermittent top-ups, which is somewhat counterintuitive as one would expect the proximity of agents to encourage any kind of interaction. It is also somewhat counterintuitive that fewer but larger withdrawals happen – we would have expected the opposite given the proximity and convenience of more granular interaction through agents.*

---

[3] Note that it is possible to have fewer deposits even though the median account has only one deposit according to the motif profile statistics because there is a non-trivial number of accounts where multiple deposits fuel the Slow Drawdown behavior.

Fast Drawdowns have a similar profile to Slow Drawdowns, except that their deposit amounts do not increase. Thus, it seems *that these users take an even larger chunk out of their initial deposit when using agents, as there is no matching increase in deposits*. The tendency for a higher average balance might suggest though that this first, large withdrawal might be happening later down the road. At any rate, a larger withdrawal is not conducive to maintaining a residual balance.

Dump-and-pulls show a clear decrease in number of deposits and in the amount of withdrawals, and also a tendency towards a lower number of withdrawals and deposit amounts. *This overall reduction in engagement across both deposits and withdrawals seem to suggest that agent usage is not conducive to these short term deposit-withdrawal cycles*. One possible reason is that the inflow of funds occurs as salaries, remittances or other incoming transfers, which may require presence at a bank to withdraw, though we do not have a way to test that hypothesis given the data we have. The other is that there are fees associated with withdrawing at agents, and withdrawing small amounts at high frequencies can quickly become prohibitively expensive. We note the occurrence of high frequencies assuming consecutive Dump-and-Pull cycles in relatively short periods of time, as there is only one withdrawal in any given Dump-and-Pull motif cycle.

The changes to average balance and number of withdrawals that we did *not* see at statistically significant levels are also quite telling. Average balances do not change, suggesting that individuals intermediate relatively similar amounts of funds, irrespective of agent usage. On one hand, this should allay some of the concerns industry observers had about agents causing a drain in saved funds by making them too easily accessible. On the other, agents do not seem to encourage a net increase in saving either – something agent banking proponents were hoping would happen due to proximity, convenience, familiarity of operator, etc. The steady number of withdrawals suggests that individuals have a set pattern for withdrawing funds, and do not deviate from that despite using agents. Again, this should help address concerns that agents make withdrawing too easy.

Changes in behavior therefore seem to be limited to the frequency of deposits and amounts of withdrawals. In so far as withdrawal amounts are driven by deposit frequencies, given that balances, withdrawal frequencies and deposit amounts do not change, we can distill the impact of agent usage to the observation that **the primary change in account usage behavior is effected through changes in deposit frequency**.

Implications of Findings

We evaluate the implications of discovering these five motifs and the impact associated with using agents by attempting to answer three questions:

- Do no-frills savings accounts further the cause of financial inclusion?
- Does agent usage encourage "desirable" behavior?
- Can these findings tell us anything about the feasibility of offering no-frills accounts?

No-Frills Accounts and Financial Inclusion

We started out this thesis by assuming that the no-frills nature of the Ordinary Savings account would allow for freeform behavior to express itself. That five behavioral motifs have emerged with distinct transactional and balance patterns is testament to the validity of that assumption. To the extent that we take specific targeting of this product and largely low-value transactional behavior in the accounts as circumstantial evidence of its use by low-income individuals, these five motifs are also proof that such individuals are capable of using formal bank accounts to meet a variety of financial intermediation needs.

Not all motifs are equally prevalent, and interestingly, the most prevalent, Dump-and-Pull at 40%, is not a savings behavior at all. However, it does point to a specific financial need being fulfilled, where individuals have a need to deposit relatively small amounts of funds for a rather short amount of time, ~US$30 for 4 days on average in this case. This may represent a salary deposit or a case where funds need to be safeguarded for a few days. It is fortuitous that no-frills accounts exist which do not charge fees on a per transaction basis that allows this behavior to exist. While not savings, this does fulfill an important financial intermediation need – one that is not satisfied by other financial instruments at their disposal, such as various commodities, credit and savings associations, informal savings collectors, money guards etc. that we explored in the section *Why Do The Poor Save?*.

The motif that most represents what we generally understand to be savings, Accumulators, is not as prevalent (13%), but nevertheless is encouraging because of the discipline these individuals display to make this possible. In the section *Why Don't The Poor Save?*, we found that it is extremely difficult for low-income individuals to save in a savings account because it is expensive, inconvenient, unreliable, seen as suboptimal use of funds compared to investing, or inadequate compared to community support during an emergency. Yet here we have clear evidence to the contrary – Bank A's Ordinary Accounts allow savings in a manner that is considered to be highly improbable in anecdotal literature.

This accretive saving behavior most closely resembles what happens in savings groups, where periodic savings result in a lump sum that is them withdrawn, making a comparison between the two useful. The duration of a typical Accumulator cycle is less than a month, while savings groups can last up to a year, implying that the discipline cannot be enforced to forego opportunities or resist temptation for as long individually as it is possible through group effort. This may not necessarily be a negative outcome, however. Group members complain of the lack of privacy, theft of a period's collection or even the entire savings pot, and inflexibility on when savings are returned. By allowing savers to save up for a month in a private, secure location, these accounts may very well be allowing payments of recurring obligations in a much more frequent manner than possible with savings groups.

The existence of these motifs is proof that FIs are successfully addressing multiple difficult challenges of financial inclusion of a market segment that was considered to be un-bankable. Accumulators are testament to the ability of low-income individuals to save with discipline and regularity in a formal instrument to arrive at a large lump sum. Slow Drawdowns demonstrate that bank accounts can be used to regularly service expense or debt obligations instead of having to resort to the less reliable money lenders or money guard or saving under the mattress. And Fast Drawdowns signal a desire to save even minimal amounts as a residual using these accounts. Thus, the findings of this thesis should be encouraging to Bank A and other FIs like it that are dedicated to offering accessible financial services for low-income individuals, as it demonstrates that no-frills savings accounts may service many different types of financial intermediation.

A natural progression from client segmentation is to offer products that are particularly tailored to a behavioral prototype, incentivizing action that is deemed “desirable.” Thus, one could argue that term deposit accounts with duration of one month may be an appropriate product to offer clients who are Accumulators and Sustained Balances, as they have demonstrated the ability to maintain a balance for a period of time. This, however, does not imply that all motifs should be mapped to new account types. A purely transactional account for the Dump-and-Pull behavior may only be cosmetically different from an Ordinary Savings account as both allow freeform deposits and withdrawals. More importantly, clients can choose to change their behavior between segments, displaying multiple motifs over a period of time, making accounts that provide the option to do that without much interference desirable.

Incidentally, these no-frills accounts have a better track record than other bank accounts even when motifs are not considered when it comes to active usage of accounts. In the section *Splicing Segments*, we found that of the 70,994 account that formed our sample and represented 1% of the Ordinary Savings Account portfolio of the FI, 41,365 (58%) of them were not dormant, in that they were undertaking transactions and their balance levels were changing. This is in contrast to the 2012 study by Dupas et. al. that found that only 18% of account holders who opened a basic savings account actively used it (Dupas, et al. 2012).

Agent Usage and Desirable Behavior

We have made a deliberate point throughout this thesis to not label any particular motif behavior as “good” or “bad” to keep the process of motif discovery as objective as possible. We temporarily set that motif-agnostic stance aside for this section to consider the impact of changes associated with agent usage. From a pro-financial inclusion perspective, anything that gets the poor to save more or interact with beneficial financial products more is desirable, as the they struggle with savings using instruments available in the informal sector.

Given that these are “savings” accounts, an increase in balances associated with agent usage would have been a desirable outcome. We see no statistically significant evidence of that. We are relieved to see that we do not see the opposite effect either, where balances decrease when agents are used. This finding should provide solace to policy makers and practitioners who are worried about the detrimental impact of agents to savings accumulation – concerns we looked into in the section *Savings and Agent Banking*. No change in balances suggests that while agent usage does not encourage individuals to bring more funds into the banking system by diverting existing sources of income or liquidating other assets, it also does not make the system leakier by encouraging client to withdraw funds, even though the option to do that is always present.

All other things being equal, we would like to see individuals having more granular control over their accounts as it allows greater flexibility in how individual funds are managed. Agent usage seems to allow just that, as their use is associated with greater engagement of Accumulators and Sustained Balances through higher deposits and withdrawals. Accumulators strive to build up a balance over time while Sustained Balances attempt to maintain a balance for a period of time. Of the five motifs, these two are the most focused on saving maintenance. It is therefore pleasing to see that greater levels of engagement occur precisely for those accounts which display such strong savings maintenance behavior. Policy makers and practitioners on the fence regarding the utility of agents can take this to evidence that at least for these behavioral types, agent usage is associated with furthering one of the goals of financial inclusion – greater interaction with formal financial instruments.

It is not a completely displeasing outcome to note that agent usage is associated with reduced Dump-and-Pull behavior. While it performs a potentially important function as a pass-through for funds, it does not result in savings accumulation and can potentially cost FIs much to support it, given the lack of float and the very short cycle times that cause many deposits and withdrawals happening in a given period of time. To the extent that we would like to see maintenance of savings balances, a reduced number of such motifs is therefore not an undesirable outcome of greater agent usage.

Feasibility of No-Frills Accounts

We end this section by looking at the business case implications of the findings of this thesis. It is well and good that FIs like Bank A are able to offer no-frills accounts that seem to promote financial inclusion, but as a for-profit bank, it needs to justify this service with respect to its bottom line. We are also in a good position to inform FIs in terms of the relative feasibility of supporting each behavioral pattern based on the existence of five distinct motifs.

Creating a complete business case is a drawn out affair that involves utilizing balance sheets and income statements, portfolio-wide transaction and balance data at various levels of segmentation, and various time allocation exercises. That manner of proprietary data is not available to us. However, we do have a close proxy in the form of stylized figures from a GAFIS Focus Note that lays out the framework for ascertaining the business case for savings accounts (BFA 2012). Bank A is comparable to one of the four FIs that were part of the GAFIS project and the stylized figures we will use were based on the study of its activities.

Conceptually speaking, there are three key elements to a business case – net interest income, fixed costs and transaction activity contribution (Table 34, reproduced verbatim from Table 2 in (BFA 2012)).

Business Case Elements
Net Interest Income (Float Income)
1. Float Revenue from Internal Treasury
2. Interest Expense Paid to Clients
Fixed Costs
3. Origination Costs (Allocated, Amortized)
4. Monthly Account Maintenance Cost (Allocated)
Transaction Activity Contribution
5. Revenue: Fee-generation Transactions
6. Expense: Transaction Costs (direct and indirect)

Table 34. Three Key Elements of a Savings Account Business Case

The net interest income depends on the average balance of the account and the duration it is held for. Float income is received from the Treasury Department of the FI, and interest is paid out to clients based on those balances. The fixed costs consist of: a) the initial acquisition cost of the account, such as issuing check books, debit cards, promotional costs etc., and b) the monthly maintenance costs that consists of allocations of various direct and indirect costs, such as salary, office rent, overhead, etc. The transaction activity contribution costs are the net of any fees charged for usage of the account, and the expenses to support such transactions, including cash handling costs, ATM maintenance, and other direct costs related to servicing and processing transactions. We can explore the impact on net interest income because we have balance data, and on transaction activity contribution because we have transaction data. We cannot explore the impact on fixed costs because we do not have relevant information.

We reproduce the summary data for average balance, number of deposit and number of withdrawals from Figure 15 in Table 35 below to help us ascertain the relative strengths of business cases for each of the motifs.

Outcome Variable	M1: Sustained Balance	M2: Dump-and-Pull	M3: Fast Drawdown	M4: Accumulator	M5: Slow Drawdown
Average Balance	KES 7,982 ($88.69)	KES 3,112 ($34.58)	KES 1,619 ($17.99)	KES 2,019 ($22.44)	KES 6,176 ($68.62)
Number of Deposits	2	1	1	2	1
Number of Withdrawals	4	1	3	2	5
Days Spanned	22	4	19	29	20

Table 35. Summary balance and transaction frequency data by motif.

Sustained Balance accounts have the best float proposition, with US$89 as the average segment balance, while Fast Drawdowns have the worst float proposition, with US$18 as the average segment balance. Despite their focus on accretive savings, Accumulators have a rather weak float proposition, with a relatively low US$22 in average balance.

In terms of transaction activity contribution, there is essentially no revenue component as no-frills accounts do not charge for basic services such as deposits or withdrawals. Every single transaction however does incur the FI a cost. The total number of deposits and withdrawals taking place can therefore serve as a proxy for this expense. At face value, it would seem that Sustained Balances and Slow Drawdowns have the most transaction related expenses, with six transactions per segment, and Dump-and-Pulls would seem to have the lowest such expenses, given only two such transactions per segment.

The contribution margin towards the bottom line will depend on the net of the float proposition and the expenses related to transactions. Thankfully the GAFIS Focus Note also provides us with stylized figures to plug in for both balance and transaction metrics, preventing us from having to make a conceptually apples-to-oranges comparison between balances and transactions. It suggests that the internal Treasury rate is 5%, the interest paid out to clients is 0.75% per annum, and the cost of a typical deposit or withdrawal transaction is US$0.72 (BFA 2012). We juxtapose these figures with those in Table 35 to arrive at indicative net contribution figures in Table 36 below. Note that since we do not have the exact figures for Bank A, these results should be considered informed estimates only.

Business Case Element	**M1: Sustained Balance**	**M2: Dump-and-Pull**	**M3: Fast Drawdown**	**M4: Accumulator**	**M5: Slow Drawdown**
Float Revenue (A)	$0.27	$0.02	$0.05	$0.09	$0.19
Interest Expense (B)	$0.04	$0.00	$0.01	$0.01	$0.03
Net Interest Income (C = A – B)	$0.23	$0.02	$0.04	$0.08	$0.16
Fee Revenue (D)	$0.00	$0.00	$0.00	$0.00	$0.00
Transaction Support Costs (E)	$4.32	$1.44	$2.88	$2.88	$4.32
Transaction Activity Net (F = D – E)	$(4.32)	$(1.44)	$(2.88)	$(2.88)	$(4.32)
Contribution Margin (C – F)	**$(4.09)**	**$(1.42)**	**$(2.84)**	**$(2.80)**	**$(4.16)**

Table 36. Contribution Margin by Motif

Thus, it seems that in all cases, the net interest income is insufficient to cover the transaction activity contribution for all motifs. This is consistent with the findings of the GAFIS Focus Note, which calculates a net loss per month, per account, of $(2.30) (BFA 2012, 6). Note however that while that $(2.30) is calculated on a per month basis, our segments are not a month long – they span from 4 days in length for Dump-and-Pulls to 29 days for Accumulators. To bring some parity to this comparison, we normalize the figures in Table 36 to see what they look like on a per 30-day basis, and present them in Table 37 below. Having a negative contribution margin may make it seem like a senseless proposition to support this savings product at all. The FIs more than make up for this by on-lending mobilized savings as retail and commercial credit at rates that are several multiples of what Treasury pays for it.

Contribution Margin	M1: Sustained Balance	M2: Dump-and-Pull	M3: Fast Drawdown	M4: Accumulator	M5: Slow Drawdown
Raw	$(4.09)	$(1.42)	$(2.84)	$(2.80)	$(4.16)
30-Day Basis	$(5.58)	$(10.68)	$(4.48)	$(2.90)	$(6.24)

Table 37. Raw and Normalized Contribution Margin, by Motif

Table 37 is not without its flaws either, however. Dump-and-Pulls are an extreme case in point, where extrapolation from a 4-day cycle to a 30-day window implies 7.5 cycles occurring within a given month. This is highly unlikely. The true monthly contribution margin is then somewhere between the raw and adjusted figures.

With all these caveats in mind, we can now infer implications to the FI's bottom line. Accumulators have the least negative impact of all four motifs. Comparatively speaking, it is therefore less costly for FIs to encourage this manner of accretive saving that is also a highly desirable exercise in terms of savings accumulation from the business case point of view. Sustained Balances and Slow Drawdowns seem to be the more expensive of the motifs, despite their higher balances. Where Dump-and-Pulls fall depends on how often they occur. If this short-term motif occurs once or twice a month, it is comparatively inexpensive to support. If however it occurs once every week, it becomes the most expensive of the motifs.

The only discernible impact associated with agent usage is that they raise deposit counts for Accumulators and Sustained Balances, while lowering it for the other three motifs. While increased deposit counts would translate to additional expenses, transacting through agent channels costs a fraction of what it costs to transact at a branch. While we do not track exactly how many of the deposits or withdrawals have occurred within a segment at an agent, we know that at least one of the deposits or withdrawals will have to have occurred through that channel. Thus, we expect the business case to actually improve as more transactions shift to the agent channel.

Further Research

We propose three particular areas for additional research. First, it is necessary to externally validate the findings of this thesis with other institutions in other countries where there are low-income populations being served by no-frills savings accounts. One reason for this is to test the validity of this technique and ensure that discernible motifs can be seen in other settings. Another reason is that Kenyans are quite familiar with engaging with FIs and have financial lives that may differ from other countries, as laid out in sections *Why Do The Poor Save?* and *Agent Banking in Kenya and Beyond*. For example, the average Kenyan low income household has ten sources of income, keeps 129% of financial income in the form of assets, and experiences 55% volatility in total monthly income. How individuals engage with low-frills accounts in settings where income amounts, sources and volatility thereof are different may impact what behavioral motifs are expressed in these accounts. Indeed, it is possible that some of these motifs may not be seen at all, while other unseen motifs may materialize.

Second, it would be immensely helpful to determine the *why* behind the motifs. In many cases, we can often surmise at some of the motivations behind the expression of the five motifs, based on previously documented anecdotal evidence combined with reasonable guesstimates. For example, Accumulators are clearly saving up towards some lump-sum for a sizeable expense, and Slow Drawdowns are seemingly using the initial deposit to pay for expenses at regular intervals. However, we don't know if Accumulators are saving up for school expenses, or for incremental home improvements, and if Fast Drawdowns leave a residual balance because it is just happened to be left behind, or if it is a deliberate attempt to save some amount of money, irrespective of how small. And we also don't know why both deposit and withdrawal frequencies decrease with agent usage even though we should expect the opposite, and why all four metrics related to deposits and withdrawals decrease with agent usage. Knowing the answers to these questions will not only help us understand the financial lives of the poor better, but also help FIs design better products and policy makers enact incentives to nudge adoption of formal financial instruments.

Whereas savings typologies have historically been generated utilizing voluminous anecdotes (as seen in section *The Poor and Their Savings*), we must now reverse that process with these motifs, seeking a sufficient number of anecdotes to appreciate the *why* behind the patterns we see. One option is to conduct a large-N survey of clients who can be linked to particular behavioral motifs, and mapping motivations for account usage as stated by them with observed data thereof. Another more involved but more comprehensive option is to conduct Financial Diaries for a smaller set of clients, tracking what is seen in their accounts as part of a larger financial ecosystem that they operate in. This will help us understand not only why bank accounts are used as they are, but also how it makes sense amongst all the other sources of incomes and expenses, as well as stocks of assets and liabilities.

Finally, it would be instructive to determine if there are demographic sub-populations for whom motif expressions are different. In the section *Why Do The Poor Save?*, we saw how savings behaviors can differ by profession and gender. Farmers have different needs than those with microenterprises. Within micro-entrepreneurs, the needs of those engaging in commerce differ from those in production or service. Women save differently than men, benefits of savings accrue differently to men and women, and sometimes, women save without men in the household even knowing about it. Accounting for demographics at the level of every account would be ideal, but collecting such data at the scale at which this study was conducted may be too difficult unless that relevant information is collected by the FI already. In such cases, large-N studies would suffice to provide a reasonable proxy.

Appendix A. Calinski–Harabasz pseudo-F Index Scores

The Calinski–Harabasz pseudo-F index score is calculated for cluster counts of three to ten, and run thirty times each, to obtain the optimal scores.

Run #	Number of Clusters							
	3	4	5	6	7	8	9	10
1	153,879.2	157,269.6	195,727.5	170,385.3	152,525.4	179,860.3	145,375.2	168,014.1
2	229,765.5	202,418.6	187,420.8	194,294.7	166,039.9	182,424.7	164,125.5	134,893.3
3	229,765.7	157,269.6	79,308.3	176,027.4	181,329.6	143,324.5	151,037.2	161,747.1
4	153,879.2	202,418.6	195,726.0	115,383.7	155,080.1	110,894.1	142,756.4	160,581.6
5	191,738.3	186,701.8	171,559.4	176,027.4	181,332.1	168,235.8	180,700.2	85,456.1
6	229,765.5	186,701.8	187,420.8	170,385.3	185,009.1	168,224.6	162,115.5	135,458.6
7	191,738.3	202,417.3	187,407.1	163,821.5	166,039.9	125,486.5	162,256.9	145,279.3
8	191,738.3	157,269.6	195,727.5	115,383.7	166,039.7	125,486.2	151,248.7	143,002.0
9	153,879.2	157,269.6	187,420.8	97,198.4	166,039.7	125,486.6	164,125.5	129,874.7
10	229,765.7	186,701.8	143,133.0	151,592.0	145,398.0	76,117.2	168,767.4	145,279.3
11	229,765.7	202,414.2	186,107.9	176,585.6	127,953.4	125,486.6	169,061.7	136,630.8
12	229,765.7	104,891.9	195,727.5	174,210.3	166,250.3	161,945.0	152,610.8	163,815.4
13	191,737.4	202,417.3	143,133.0	170,385.3	184,116.4	151,549.0	164,125.5	112,361.1
14	153,879.2	104,891.9	195,726.0	151,592.0	166,039.9	161,945.9	143,914.4	127,440.7
15	229,765.7	203,831.8	173,697.3	174,210.3	156,899.1	143,324.7	142,756.1	135,458.6
16	191,737.4	186,701.8	185,231.4	176,585.8	166,039.6	110,895.0	145,375.3	129,874.7
17	229,765.7	202,417.5	187,420.8	153,380.7	166,609.5	182,425.1	162,827.0	145,616.4
18	229,765.5	186,701.8	187,420.8	170,385.3	184,116.4	143,321.1	164,516.3	127,502.5
19	229,765.5	157,269.6	187,420.8	174,210.0	165,847.8	161,986.1	169,715.0	128,580.7
20	229,765.5	203,847.5	185,231.4	194,295.2	184,283.3	163,257.2	126,002.2	144,470.1
21	229,765.7	186,701.8	195,727.5	194,295.2	166,039.7	168,235.8	162,255.9	105,360.4
22	153,879.2	202,537.6	79,308.3	152,416.8	96,624.8	141,718.3	159,813.4	135,439.5
23	153,879.2	203,831.8	187,420.8	194,295.2	184,283.3	110,894.1	140,615.2	143,002.0
24	229,765.7	203,832.4	187,420.8	194,294.7	148,707.6	161,986.1	162,827.0	160,669.7
25	229,765.7	203,852.2	185,231.4	151,592.0	166,039.3	159,143.6	126,032.4	161,747.1
26	153,879.2	157,269.6	185,231.4	176,027.5	181,332.6	159,627.5	164,124.0	149,271.6
27	153,879.2	157,269.6	185,231.1	194,295.2	185,377.3	110,232.2	159,347.1	175,192.7
28	153,879.2	202,537.6	186,107.9	194,294.7	127,953.4	182,424.7	151,273.2	133,906.6
29	191,737.4	202,417.4	186,107.9	194,295.2	184,283.3	143,321.1	126,002.2	125,108.0
30	229,765.7	202,417.3	171,559.4	151,591.5	184,116.4	181,194.3	126,002.2	127,502.5
Maximum	229,765.7	203,852.2	195,727.5	194,295.2	185,377.3	182,425.1	180,700.2	175,192.7
Mean	199,394.1	182,483.0	176,910.5	168,124.6	165,258.2	147,681.8	153,723.5	139,284.6
Median	210,751.9	194,558.0	186,757.5	174,210.2	166,039.9	155,346.3	159,580.2	136,044.7
Minimum	153,879.2	104,891.9	79,308.3	97,198.4	96,624.8	76,117.2	126,002.2	85,456.1

Appendix B. Detailed Results of AB-GMM Runs

```
.
. foreach this_var of varlist balance_log dep_cnt_log wd_cnt_log dep_avg_log wd_avg_log {
  2.
.         di _n _n _n "********** `this_var' **********" _n
  3.
.
.         di _n "---------- `this_var': AGENTs & MOTIFs, w INTERACTION ----------" _n
  4.                xi:  xtabond2  `this_var'  L.`this_var'  i.used_agent  i.motif  L.i.motif
i.used_agent*i.motif, ///
>                gmm(`this_var', lag(2 2)) gmm(i.motif, lag(2 2) eq(level)) iv(i.used_agent,
eq(level)) ///
>                twostep robust ortho nodiffsargan
  5.         test (_Imotif_2) == (_Imotif_3) == (_Imotif_4) == (_Imotif_5)
  6.         test (_Imotif_2) == (_Imotif_3) == (_Imotif_4) == (_Imotif_5) == 0
  7.         test (L1._Imotif_2) == (L1._Imotif_3) == (L1._Imotif_4) == (L1._Imotif_5)
  8.         test (L1._Imotif_2) == (L1._Imotif_3) == (L1._Imotif_4) == (L1._Imotif_5) == 0
  9.         test (_IuseXmot_1_2) == (_IuseXmot_1_3) == (_IuseXmot_1_4) == (_IuseXmot_1_5)
 10.         test (_IuseXmot_1_2) == (_IuseXmot_1_3) == (_IuseXmot_1_4) == (_IuseXmot_1_5) ==
0
 11.
.         // Added after defense - pairwise tests
.         test (_Imotif_2) == (_Imotif_3)
 12.         test (_Imotif_2) == (_Imotif_4)
 13.         test (_Imotif_2) == (_Imotif_5)
 14.         test (_Imotif_3) == (_Imotif_4)
 15.         test (_Imotif_3) == (_Imotif_5)
 16.         test (_Imotif_4) == (_Imotif_5)
 17.
.         test (L1._Imotif_2) == (L1._Imotif_3)
 18.         test (L1._Imotif_2) == (L1._Imotif_4)
 19.         test (L1._Imotif_2) == (L1._Imotif_5)
 20.         test (L1._Imotif_3) == (L1._Imotif_4)
 21.         test (L1._Imotif_3) == (L1._Imotif_5)
 22.         test (L1._Imotif_4) == (L1._Imotif_5)
 23.
.         test (_IuseXmot_1_2) == (_IuseXmot_1_3)
 24.         test (_IuseXmot_1_2) == (_IuseXmot_1_4)
 25.         test (_IuseXmot_1_2) == (_IuseXmot_1_5)
 26.         test (_IuseXmot_1_3) == (_IuseXmot_1_4)
 27.         test (_IuseXmot_1_3) == (_IuseXmot_1_5)
 28.         test (_IuseXmot_1_4) == (_IuseXmot_1_5)
 29. }
```

```
********** balance_log **********

i.used_agent      _Iused_agen_0-1     (naturally coded; _Iused_agen_0 omitted)
i.motif           _Imotif_1-5         (naturally coded; _Imotif_1 omitted)
i.u~ent*i.motif   _IuseXmot_#_#       (coded as above)
Favoring space over speed. To switch, type or click on mata: mata set matafavor speed, perm.
_Iused_agen_1 dropped due to collinearity
_Imotif_2 dropped due to collinearity
_Imotif_3 dropped due to collinearity
_Imotif_4 dropped due to collinearity
_Imotif_5 dropped due to collinearity
Warning: Two-step estimated covariance matrix of moments is singular.
  Using a generalized inverse to calculate optimal weighting matrix for two-step estimation.

Dynamic panel-data estimation, two-step system GMM
------------------------------------------------------------------------------
Group variable: acctid_n                        Number of obs      =    228481
Time variable : fake_year                       Number of groups   =     23097
Number of instruments = 196                     Obs per group: min =         1
Wald chi2(14) =    315.71                                      avg =      9.89
Prob > chi2   =     0.000                                      max =        34
------------------------------------------------------------------------------
             |              Corrected
 balance_log |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
 balance_log |
         L1. |   -.101182   .0307069    -3.30   0.001    -.1613664   -.0409976
             |
_Iused_age~1 |  -.5188923    .764558    -0.68   0.497    -2.017398    .9796138
   _Imotif_2 |  -2.003237    .504521    -3.97   0.000     -2.99208   -1.014394
   _Imotif_3 |  -2.064502   .6219075    -3.32   0.001    -3.283418   -.8455853
   _Imotif_4 |  -3.767742   .6795154    -5.54   0.000    -5.099568   -2.435916
   _Imotif_5 |  -.7071874   .6749266    -1.05   0.295    -2.030019    .6156444
             |
   _Imotif_2 |
         L1. |  -1.813899   .3550682    -5.11   0.000     -2.50982   -1.117978
             |
   _Imotif_3 |
         L1. |  -2.455033   .4258714    -5.76   0.000    -3.289725    -1.62034
             |
   _Imotif_4 |
         L1. |  -1.012744   .7374869    -1.37   0.170    -2.458192    .4327038
             |
   _Imotif_5 |
         L1. |  -1.792886   .5273496    -3.40   0.001    -2.826472   -.7592995
             |
_IuseXmot_~2 |   .5472059   .9908995     0.55   0.581    -1.394921    2.489333
_IuseXmot_~3 |    1.45399   1.294337     1.12   0.261    -1.082864    3.990844
_IuseXmot_~4 |   .4627034   1.348873     0.34   0.732    -2.181039    3.106446
_IuseXmot_~5 |    .315488   1.064299     0.30   0.767      -1.7705    2.401475
       _cons |   11.97244   .6117569    19.57   0.000     10.77342    13.17146
------------------------------------------------------------------------------
Instruments for orthogonal deviations equation
  GMM-type (missing=0, separate instruments for each period unless collapsed)
    L2.balance_log
Instruments for levels equation
  Standard
    _Iused_agen_1
    _cons
  GMM-type (missing=0, separate instruments for each period unless collapsed)
    DL2.(1b.motif 2.motif 3.motif 4.motif 5.motif)
    DL.balance_log
------------------------------------------------------------------------------
Arellano-Bond test for AR(1) in first differences: z = -10.80  Pr > z =  0.000
Arellano-Bond test for AR(2) in first differences: z =  -0.43  Pr > z =  0.671
------------------------------------------------------------------------------
Sargan test of overid. restrictions: chi2(181)  = 337.60  Prob > chi2 =  0.000
  (Not robust, but not weakened by many instruments.)
Hansen test of overid. restrictions: chi2(181)  = 223.69  Prob > chi2 =  0.017
  (Robust, but weakened by many instruments.)
```

```
( 1)  _Imotif_2 - _Imotif_3 = 0
( 2)  _Imotif_2 - _Imotif_4 = 0
( 3)  _Imotif_2 - _Imotif_5 = 0

          chi2(  3) =   25.22
        Prob > chi2 =    0.0000

( 1)  _Imotif_2 - _Imotif_3 = 0
( 2)  _Imotif_2 - _Imotif_4 = 0
( 3)  _Imotif_2 - _Imotif_5 = 0
( 4)  _Imotif_2 = 0

          chi2(  4) =   65.36
        Prob > chi2 =    0.0000

( 1)  L._Imotif_2 - L._Imotif_3 = 0
( 2)  L._Imotif_2 - L._Imotif_4 = 0
( 3)  L._Imotif_2 - L._Imotif_5 = 0

          chi2(  3) =    5.54
        Prob > chi2 =    0.1360

( 1)  L._Imotif_2 - L._Imotif_3 = 0
( 2)  L._Imotif_2 - L._Imotif_4 = 0
( 3)  L._Imotif_2 - L._Imotif_5 = 0
( 4)  L._Imotif_2 = 0

          chi2(  4) =   53.61
        Prob > chi2 =    0.0000

( 1)  _IuseXmot_1_2 - _IuseXmot_1_3 = 0
( 2)  _IuseXmot_1_2 - _IuseXmot_1_4 = 0
( 3)  _IuseXmot_1_2 - _IuseXmot_1_5 = 0

          chi2(  3) =    0.87
        Prob > chi2 =    0.8318

( 1)  _IuseXmot_1_2 - _IuseXmot_1_3 = 0
( 2)  _IuseXmot_1_2 - _IuseXmot_1_4 = 0
( 3)  _IuseXmot_1_2 - _IuseXmot_1_5 = 0
( 4)  _IuseXmot_1_2 = 0

          chi2(  4) =    1.32
        Prob > chi2 =    0.8573

( 1)  _Imotif_2 - _Imotif_3 = 0

          chi2(  1) =    0.02
        Prob > chi2 =    0.9024

( 1)  _Imotif_2 - _Imotif_4 = 0

          chi2(  1) =    6.62
        Prob > chi2 =    0.0101

( 1)  _Imotif_2 - _Imotif_5 = 0

          chi2(  1) =    9.48
        Prob > chi2 =    0.0021

( 1)  _Imotif_3 - _Imotif_4 = 0

          chi2(  1) =    5.56
        Prob > chi2 =    0.0184

( 1)  _Imotif_3 - _Imotif_5 = 0

          chi2(  1) =    4.49
        Prob > chi2 =    0.0340
```

```
( 1)  _Imotif_4 - _Imotif_5 = 0

          chi2(  1) =   20.80
        Prob > chi2 =    0.0000

( 1)  L._Imotif_2 - L._Imotif_3 = 0

          chi2(  1) =    3.28
        Prob > chi2 =    0.0701

( 1)  L._Imotif_2 - L._Imotif_4 = 0

          chi2(  1) =    2.25
        Prob > chi2 =    0.1333

( 1)  L._Imotif_2 - L._Imotif_5 = 0

          chi2(  1) =    0.00
        Prob > chi2 =    0.9537

( 1)  L._Imotif_3 - L._Imotif_4 = 0

          chi2(  1) =    5.02
        Prob > chi2 =    0.0250

( 1)  L._Imotif_3 - L._Imotif_5 = 0

          chi2(  1) =    1.96
        Prob > chi2 =    0.1620

( 1)  L._Imotif_4 - L._Imotif_5 = 0

          chi2(  1) =    1.49
        Prob > chi2 =    0.2226

( 1)  _IuseXmot_1_2 - _IuseXmot_1_3 = 0

          chi2(  1) =    0.67
        Prob > chi2 =    0.4137

( 1)  _IuseXmot_1_2 - _IuseXmot_1_4 = 0

          chi2(  1) =    0.01
        Prob > chi2 =    0.9424

( 1)  _IuseXmot_1_2 - _IuseXmot_1_5 = 0

          chi2(  1) =    0.08
        Prob > chi2 =    0.7795

( 1)  _IuseXmot_1_3 - _IuseXmot_1_4 = 0

          chi2(  1) =    0.43
        Prob > chi2 =    0.5141

( 1)  _IuseXmot_1_3 - _IuseXmot_1_5 = 0

          chi2(  1) =    0.83
        Prob > chi2 =    0.3619

( 1)  _IuseXmot_1_4 - _IuseXmot_1_5 = 0

          chi2(  1) =    0.01
        Prob > chi2 =    0.9081
```

```
********** dep_cnt_log **********

i.used_agent      _Iused_agen_0-1     (naturally coded; _Iused_agen_0 omitted)
i.motif           _Imotif_1-5         (naturally coded; _Imotif_1 omitted)
i.u~ent*i.motif   _IuseXmot_#_#       (coded as above)
Favoring space over speed. To switch, type or click on mata: mata set matafavor speed, perm.
_Iused_agen_1 dropped due to collinearity
_Imotif_2 dropped due to collinearity
_Imotif_3 dropped due to collinearity
_Imotif_4 dropped due to collinearity
_Imotif_5 dropped due to collinearity
Warning: Two-step estimated covariance matrix of moments is singular.
  Using a generalized inverse to calculate optimal weighting matrix for two-step estimation.

Dynamic panel-data estimation, two-step system GMM
------------------------------------------------------------------------------
Group variable: acctid_n                        Number of obs      =    176132
Time variable : fake_year                       Number of groups   =     21610
Number of instruments = 196                     Obs per group: min =         1
Wald chi2(14) =    379.07                                      avg =      8.15
Prob > chi2   =     0.000                                      max =        34
------------------------------------------------------------------------------
             |              Corrected
 dep_cnt_log |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
 dep_cnt_log |
         L1. |   .1064786   .0160393     6.64   0.000     .0750422     .137915
             |
_Iused_age~1 |   1.170186   .4477236     2.61   0.009     .2926643    2.047709
   _Imotif_2 |   .2095436   .3219806     0.65   0.515    -.4215268     .840614
   _Imotif_3 |  -.9169092   .3603158    -2.54   0.011    -1.623115   -.2107032
   _Imotif_4 |  -2.143306   .3963232    -5.41   0.000    -2.920085   -1.366527
   _Imotif_5 |  -.3904199   .3827726    -1.02   0.308     -1.14064    .3598007
             |
   _Imotif_2 |
         L1. |  -.7549923   .2522347    -2.99   0.003    -1.249363   -.2606214
             |
   _Imotif_3 |
         L1. |  -.8596226   .3110202    -2.76   0.006    -1.469211   -.2500343
             |
   _Imotif_4 |
         L1. |  -.3319039   .3904649    -0.85   0.395    -1.097201    .4333932
             |
   _Imotif_5 |
         L1. |  -.3666635   .3260686    -1.12   0.261    -1.005746    .2724191
             |
_IuseXmot_~2 |   -1.69725   .5707557    -2.97   0.003     -2.81591   -.5785891
_IuseXmot_~3 |   -1.53431   .7412336    -2.07   0.038    -2.987102   -.0815193
_IuseXmot_~4 |  -.2046163    .766769    -0.27   0.790    -1.707456    1.298223
_IuseXmot_~5 |  -1.554072   .6415151    -2.42   0.015    -2.811418   -.2967253
       _cons |   2.178705   .3478832     6.26   0.000     1.496866    2.860543
------------------------------------------------------------------------------
Instruments for orthogonal deviations equation
  GMM-type (missing=0, separate instruments for each period unless collapsed)
    L2.dep_cnt_log
Instruments for levels equation
  Standard
    _Iused_agen_1
    _cons
  GMM-type (missing=0, separate instruments for each period unless collapsed)
    DL2.(1b.motif 2.motif 3.motif 4.motif 5.motif)
    DL.dep_cnt_log
------------------------------------------------------------------------------
Arellano-Bond test for AR(1) in first differences: z =  -8.26  Pr > z =  0.000
Arellano-Bond test for AR(2) in first differences: z =  -0.18  Pr > z =  0.857
------------------------------------------------------------------------------
Sargan test of overid. restrictions: chi2(181)  = 305.00  Prob > chi2 =  0.000
  (Not robust, but not weakened by many instruments.)
Hansen test of overid. restrictions: chi2(181)  = 224.19  Prob > chi2 =  0.016
  (Robust, but weakened by many instruments.)
```

```
( 1)  _Imotif_2 - _Imotif_3 = 0
( 2)  _Imotif_2 - _Imotif_4 = 0
( 3)  _Imotif_2 - _Imotif_5 = 0

          chi2(  3) =   43.91
        Prob > chi2 =    0.0000

( 1)  _Imotif_2 - _Imotif_3 = 0
( 2)  _Imotif_2 - _Imotif_4 = 0
( 3)  _Imotif_2 - _Imotif_5 = 0
( 4)  _Imotif_2 = 0

          chi2(  4) =   50.83
        Prob > chi2 =    0.0000

( 1)  L._Imotif_2 - L._Imotif_3 = 0
( 2)  L._Imotif_2 - L._Imotif_4 = 0
( 3)  L._Imotif_2 - L._Imotif_5 = 0

          chi2(  3) =    8.71
        Prob > chi2 =    0.0334

( 1)  L._Imotif_2 - L._Imotif_3 = 0
( 2)  L._Imotif_2 - L._Imotif_4 = 0
( 3)  L._Imotif_2 - L._Imotif_5 = 0
( 4)  L._Imotif_2 = 0

          chi2(  4) =   22.18
        Prob > chi2 =    0.0002

( 1)  _IuseXmot_1_2 - _IuseXmot_1_3 = 0
( 2)  _IuseXmot_1_2 - _IuseXmot_1_4 = 0
( 3)  _IuseXmot_1_2 - _IuseXmot_1_5 = 0

          chi2(  3) =    5.90
        Prob > chi2 =    0.1166

( 1)  _IuseXmot_1_2 - _IuseXmot_1_3 = 0
( 2)  _IuseXmot_1_2 - _IuseXmot_1_4 = 0
( 3)  _IuseXmot_1_2 - _IuseXmot_1_5 = 0
( 4)  _IuseXmot_1_2 = 0

          chi2(  4) =   13.71
        Prob > chi2 =    0.0083

( 1)  _Imotif_2 - _Imotif_3 = 0

          chi2(  1) =   15.40
        Prob > chi2 =    0.0001

( 1)  _Imotif_2 - _Imotif_4 = 0

          chi2(  1) =   36.48
        Prob > chi2 =    0.0000

( 1)  _Imotif_2 - _Imotif_5 = 0

          chi2(  1) =    4.65
        Prob > chi2 =    0.0311

( 1)  _Imotif_3 - _Imotif_4 = 0

          chi2(  1) =    8.02
        Prob > chi2 =    0.0046

( 1)  _Imotif_3 - _Imotif_5 = 0

          chi2(  1) =    2.10
        Prob > chi2 =    0.1473
```

```
( 1)  _Imotif_4 - _Imotif_5 = 0

          chi2(  1) =   18.52
        Prob > chi2 =    0.0000

( 1)  L._Imotif_2 - L._Imotif_3 = 0

          chi2(  1) =    0.43
        Prob > chi2 =    0.5101

( 1)  L._Imotif_2 - L._Imotif_4 = 0

          chi2(  1) =    2.09
        Prob > chi2 =    0.1484

( 1)  L._Imotif_2 - L._Imotif_5 = 0

          chi2(  1) =    5.13
        Prob > chi2 =    0.0235

( 1)  L._Imotif_3 - L._Imotif_4 = 0

          chi2(  1) =    2.15
        Prob > chi2 =    0.1425

( 1)  L._Imotif_3 - L._Imotif_5 = 0

          chi2(  1) =    3.20
        Prob > chi2 =    0.0734

( 1)  L._Imotif_4 - L._Imotif_5 = 0

          chi2(  1) =    0.01
        Prob > chi2 =    0.9225

( 1)  _IuseXmot_1_2 - _IuseXmot_1_3 = 0

          chi2(  1) =    0.07
        Prob > chi2 =    0.7929

( 1)  _IuseXmot_1_2 - _IuseXmot_1_4 = 0

          chi2(  1) =    5.74
        Prob > chi2 =    0.0166

( 1)  _IuseXmot_1_2 - _IuseXmot_1_5 = 0

          chi2(  1) =    0.08
        Prob > chi2 =    0.7743

( 1)  _IuseXmot_1_3 - _IuseXmot_1_4 = 0

          chi2(  1) =    2.58
        Prob > chi2 =    0.1079

( 1)  _IuseXmot_1_3 - _IuseXmot_1_5 = 0

          chi2(  1) =    0.00
        Prob > chi2 =    0.9764

( 1)  _IuseXmot_1_4 - _IuseXmot_1_5 = 0

          chi2(  1) =    3.90
        Prob > chi2 =    0.0484
```

```
********** wd_cnt_log **********

i.used_agent      _Iused_agen_0-1     (naturally coded; _Iused_agen_0 omitted)
i.motif           _Imotif_1-5         (naturally coded; _Imotif_1 omitted)
i.u~ent*i.motif   _IuseXmot_#_#       (coded as above)
Favoring space over speed. To switch, type or click on mata: mata set matafavor speed, perm.
_Iused_agen_1 dropped due to collinearity
_Imotif_2 dropped due to collinearity
_Imotif_3 dropped due to collinearity
_Imotif_4 dropped due to collinearity
_Imotif_5 dropped due to collinearity
Warning: Two-step estimated covariance matrix of moments is singular.
  Using a generalized inverse to calculate optimal weighting matrix for two-step estimation.

Dynamic panel-data estimation, two-step system GMM
------------------------------------------------------------------------------
Group variable: acctid_n                        Number of obs      =    217826
Time variable : fake_year                       Number of groups   =     22606
Number of instruments = 196                     Obs per group: min =         1
Wald chi2(14) =    367.94                                      avg =      9.64
Prob > chi2   =     0.000                                      max =        34
------------------------------------------------------------------------------
             |              Corrected
  wd_cnt_log |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
  wd_cnt_log |
         L1. |   .1131244   .0296804     3.81   0.000     .0549518     .171297
             |
_Iused_age~1 |   .3534793   .5116903     0.69   0.490    -.6494153    1.356374
   _Imotif_2 |  -1.583292   .3383066    -4.68   0.000    -2.246361   -.9202234
   _Imotif_3 |  -1.156415   .4252267    -2.72   0.007    -1.989844   -.3229856
   _Imotif_4 |  -3.080628   .5185109    -5.94   0.000     -4.09689   -2.064365
   _Imotif_5 |  -.1979846   .4287083    -0.46   0.644    -1.038237    .6422681
             |
   _Imotif_2 |
         L1. |  -.1343781   .2376121    -0.57   0.572    -.6000893    .3313331
             |
   _Imotif_3 |
         L1. |   .0594082   .3474252     0.17   0.864    -.6215327    .7403492
             |
   _Imotif_4 |
         L1. |   .2856343     .45225     0.63   0.528    -.6007593    1.172028
             |
   _Imotif_5 |
         L1. |   .4793437   .3525221     1.36   0.174    -.2115869    1.170274
             |
_IuseXmot_~2 |  -.6557262   .6559942    -1.00   0.318    -1.941451    .6299989
_IuseXmot_~3 |  -.6167269   .8492438    -0.73   0.468    -2.281214     1.04776
_IuseXmot_~4 |   .3934478    .930022     0.42   0.672    -1.429362    2.216257
_IuseXmot_~5 |  -.6503641   .7149614    -0.91   0.363    -2.051663    .7509344
       _cons |   3.006895   .3802394     7.91   0.000      2.26164    3.752151
------------------------------------------------------------------------------
Instruments for orthogonal deviations equation
  GMM-type (missing=0, separate instruments for each period unless collapsed)
    L2.wd_cnt_log
Instruments for levels equation
  Standard
    _Iused_agen_1
    _cons
  GMM-type (missing=0, separate instruments for each period unless collapsed)
    DL2.(1b.motif 2.motif 3.motif 4.motif 5.motif)
    DL.wd_cnt_log
------------------------------------------------------------------------------
Arellano-Bond test for AR(1) in first differences: z =  -8.40  Pr > z =  0.000
Arellano-Bond test for AR(2) in first differences: z =  -0.24  Pr > z =  0.808
------------------------------------------------------------------------------
Sargan test of overid. restrictions: chi2(181)  = 288.65  Prob > chi2 =  0.000
  (Not robust, but not weakened by many instruments.)
Hansen test of overid. restrictions: chi2(181)  = 238.84  Prob > chi2 =  0.003
  (Robust, but weakened by many instruments.)
```

```
( 1)  _Imotif_2 - _Imotif_3 = 0
( 2)  _Imotif_2 - _Imotif_4 = 0
( 3)  _Imotif_2 - _Imotif_5 = 0

          chi2(  3) =   65.25
        Prob > chi2 =    0.0000

( 1)  _Imotif_2 - _Imotif_3 = 0
( 2)  _Imotif_2 - _Imotif_4 = 0
( 3)  _Imotif_2 - _Imotif_5 = 0
( 4)  _Imotif_2 = 0

          chi2(  4) =   99.52
        Prob > chi2 =    0.0000

( 1)  L._Imotif_2 - L._Imotif_3 = 0
( 2)  L._Imotif_2 - L._Imotif_4 = 0
( 3)  L._Imotif_2 - L._Imotif_5 = 0

          chi2(  3) =   13.63
        Prob > chi2 =    0.0035

( 1)  L._Imotif_2 - L._Imotif_3 = 0
( 2)  L._Imotif_2 - L._Imotif_4 = 0
( 3)  L._Imotif_2 - L._Imotif_5 = 0
( 4)  L._Imotif_2 = 0

          chi2(  4) =   17.83
        Prob > chi2 =    0.0013

( 1)  _IuseXmot_1_2 - _IuseXmot_1_3 = 0
( 2)  _IuseXmot_1_2 - _IuseXmot_1_4 = 0
( 3)  _IuseXmot_1_2 - _IuseXmot_1_5 = 0

          chi2(  3) =    2.05
        Prob > chi2 =    0.5628

( 1)  _IuseXmot_1_2 - _IuseXmot_1_3 = 0
( 2)  _IuseXmot_1_2 - _IuseXmot_1_4 = 0
( 3)  _IuseXmot_1_2 - _IuseXmot_1_5 = 0
( 4)  _IuseXmot_1_2 = 0

          chi2(  4) =    3.00
        Prob > chi2 =    0.5571

( 1)  _Imotif_2 - _Imotif_3 = 0

          chi2(  1) =    1.76
        Prob > chi2 =    0.1847

( 1)  _Imotif_2 - _Imotif_4 = 0

          chi2(  1) =    9.33
        Prob > chi2 =    0.0022

( 1)  _Imotif_2 - _Imotif_5 = 0

          chi2(  1) =   25.20
        Prob > chi2 =    0.0000

( 1)  _Imotif_3 - _Imotif_4 = 0

          chi2(  1) =   13.31
        Prob > chi2 =    0.0003

( 1)  _Imotif_3 - _Imotif_5 = 0

          chi2(  1) =    4.79
        Prob > chi2 =    0.0287
```

```
( 1)  _Imotif_4 - _Imotif_5 = 0

          chi2(  1) =   39.59
        Prob > chi2 =    0.0000

( 1)  L._Imotif_2 - L._Imotif_3 = 0

          chi2(  1) =    0.69
        Prob > chi2 =    0.4063

( 1)  L._Imotif_2 - L._Imotif_4 = 0

          chi2(  1) =    1.27
        Prob > chi2 =    0.2605

( 1)  L._Imotif_2 - L._Imotif_5 = 0

          chi2(  1) =   10.14
        Prob > chi2 =    0.0014

( 1)  L._Imotif_3 - L._Imotif_4 = 0

          chi2(  1) =    0.26
        Prob > chi2 =    0.6069

( 1)  L._Imotif_3 - L._Imotif_5 = 0

          chi2(  1) =    1.70
        Prob > chi2 =    0.1925

( 1)  L._Imotif_4 - L._Imotif_5 = 0

          chi2(  1) =    0.20
        Prob > chi2 =    0.6538

( 1)  _IuseXmot_1_2 - _IuseXmot_1_3 = 0

          chi2(  1) =    0.00
        Prob > chi2 =    0.9572

( 1)  _IuseXmot_1_2 - _IuseXmot_1_4 = 0

          chi2(  1) =    1.76
        Prob > chi2 =    0.1850

( 1)  _IuseXmot_1_2 - _IuseXmot_1_5 = 0

          chi2(  1) =    0.00
        Prob > chi2 =    0.9925

( 1)  _IuseXmot_1_3 - _IuseXmot_1_4 = 0

          chi2(  1) =    1.01
        Prob > chi2 =    0.3147

( 1)  _IuseXmot_1_3 - _IuseXmot_1_5 = 0

          chi2(  1) =    0.00
        Prob > chi2 =    0.9672

( 1)  _IuseXmot_1_4 - _IuseXmot_1_5 = 0

          chi2(  1) =    1.80
        Prob > chi2 =    0.1802
```

```
********** dep_avg_log **********

i.used_agent      _Iused_agen_0-1     (naturally coded; _Iused_agen_0 omitted)
i.motif           _Imotif_1-5         (naturally coded; _Imotif_1 omitted)
i.u~ent*i.motif   _IuseXmot_#_#       (coded as above)
Favoring space over speed. To switch, type or click on mata: mata set matafavor speed, perm.
_Iused_agen_1 dropped due to collinearity
_Imotif_2 dropped due to collinearity
_Imotif_3 dropped due to collinearity
_Imotif_4 dropped due to collinearity
_Imotif_5 dropped due to collinearity
Warning: Two-step estimated covariance matrix of moments is singular.
  Using a generalized inverse to calculate optimal weighting matrix for two-step estimation.

Dynamic panel-data estimation, two-step system GMM
------------------------------------------------------------------------------
Group variable: acctid_n                        Number of obs      =    176132
Time variable : fake_year                       Number of groups   =     21610
Number of instruments = 196                     Obs per group: min =         1
Wald chi2(14) =    485.34                                      avg =      8.15
Prob > chi2   =     0.000                                      max =        34
------------------------------------------------------------------------------
             |              Corrected
 dep_avg_log |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
 dep_avg_log |
         L1. |   .0991747    .025187     3.94   0.000     .0498091    .1485403
             |
_Iused_age~1 |  -.6420674   .5241197    -1.23   0.221    -1.669323    .3851884
   _Imotif_2 |   .9902483   .3890282     2.55   0.011     .2277671     1.75273
   _Imotif_3 |   .2515028   .4626003     0.54   0.587    -.6551772    1.158183
   _Imotif_4 |   2.510915   .4521045     5.55   0.000     1.624807    3.397024
   _Imotif_5 |   .4090586   .4542761     0.90   0.368    -.4813063    1.299423
             |
   _Imotif_2 |
         L1. |   .9346831   .2731379     3.42   0.001     .3993426    1.470024
             |
   _Imotif_3 |
         L1. |   .5323351   .3257472     1.63   0.102    -.1061177    1.170788
             |
   _Imotif_4 |
         L1. |   1.135126   .4314301     2.63   0.009     .2895382    1.980713
             |
   _Imotif_5 |
         L1. |   .2694253   .3707872     0.73   0.467    -.4573043    .9961548
             |
_IuseXmot_~2 |  -.4370641   .6892106    -0.63   0.526    -1.787892    .9137638
_IuseXmot_~3 |    .638654   .9368172     0.68   0.495    -1.197474    2.474782
_IuseXmot_~4 |   .3538001   .9202411     0.38   0.701    -1.449839    2.157439
_IuseXmot_~5 |   2.125975   .7425196     2.86   0.004     .6706632    3.581287
       _cons |   6.649983   .4070304    16.34   0.000     5.852218    7.447747
------------------------------------------------------------------------------
Instruments for orthogonal deviations equation
  GMM-type (missing=0, separate instruments for each period unless collapsed)
    L2.dep_avg_log
Instruments for levels equation
  Standard
    _Iused_agen_1
    _cons
  GMM-type (missing=0, separate instruments for each period unless collapsed)
    DL2.(1b.motif 2.motif 3.motif 4.motif 5.motif)
    DL.dep_avg_log
------------------------------------------------------------------------------
Arellano-Bond test for AR(1) in first differences: z =  -7.07  Pr > z =  0.000
Arellano-Bond test for AR(2) in first differences: z =  -1.09  Pr > z =  0.277
------------------------------------------------------------------------------
Sargan test of overid. restrictions: chi2(181)  = 434.70  Prob > chi2 =  0.000
  (Not robust, but not weakened by many instruments.)
Hansen test of overid. restrictions: chi2(181)  = 260.25  Prob > chi2 =  0.000
  (Robust, but weakened by many instruments.)
```

```
 ( 1)  _Imotif_2 - _Imotif_3 = 0
 ( 2)  _Imotif_2 - _Imotif_4 = 0
 ( 3)  _Imotif_2 - _Imotif_5 = 0

           chi2(  3) =   31.51
         Prob > chi2 =    0.0000

 ( 1)  _Imotif_2 - _Imotif_3 = 0
 ( 2)  _Imotif_2 - _Imotif_4 = 0
 ( 3)  _Imotif_2 - _Imotif_5 = 0
 ( 4)  _Imotif_2 = 0

           chi2(  4) =   45.24
         Prob > chi2 =    0.0000

 ( 1)  L._Imotif_2 - L._Imotif_3 = 0
 ( 2)  L._Imotif_2 - L._Imotif_4 = 0
 ( 3)  L._Imotif_2 - L._Imotif_5 = 0

           chi2(  3) =    8.46
         Prob > chi2 =    0.0373

 ( 1)  L._Imotif_2 - L._Imotif_3 = 0
 ( 2)  L._Imotif_2 - L._Imotif_4 = 0
 ( 3)  L._Imotif_2 - L._Imotif_5 = 0
 ( 4)  L._Imotif_2 = 0

           chi2(  4) =   15.72
         Prob > chi2 =    0.0034

 ( 1)  _IuseXmot_1_2 - _IuseXmot_1_3 = 0
 ( 2)  _IuseXmot_1_2 - _IuseXmot_1_4 = 0
 ( 3)  _IuseXmot_1_2 - _IuseXmot_1_5 = 0

           chi2(  3) =   18.32
         Prob > chi2 =    0.0004

 ( 1)  _IuseXmot_1_2 - _IuseXmot_1_3 = 0
 ( 2)  _IuseXmot_1_2 - _IuseXmot_1_4 = 0
 ( 3)  _IuseXmot_1_2 - _IuseXmot_1_5 = 0
 ( 4)  _IuseXmot_1_2 = 0

           chi2(  4) =   19.20
         Prob > chi2 =    0.0007

 ( 1)  _Imotif_2 - _Imotif_3 = 0

           chi2(  1) =    4.16
         Prob > chi2 =    0.0413

 ( 1)  _Imotif_2 - _Imotif_4 = 0

           chi2(  1) =   11.18
         Prob > chi2 =    0.0008

 ( 1)  _Imotif_2 - _Imotif_5 = 0

           chi2(  1) =    3.01
         Prob > chi2 =    0.0827

 ( 1)  _Imotif_3 - _Imotif_4 = 0

           chi2(  1) =   20.21
         Prob > chi2 =    0.0000

 ( 1)  _Imotif_3 - _Imotif_5 = 0

           chi2(  1) =    0.10
         Prob > chi2 =    0.7500
```

```
( 1)  _Imotif_4 - _Imotif_5 = 0

          chi2(  1) =   20.39
        Prob > chi2 =    0.0000

( 1)  L._Imotif_2 - L._Imotif_3 = 0

          chi2(  1) =    2.20
        Prob > chi2 =    0.1377

( 1)  L._Imotif_2 - L._Imotif_4 = 0

          chi2(  1) =    0.50
        Prob > chi2 =    0.4817

( 1)  L._Imotif_2 - L._Imotif_5 = 0

          chi2(  1) =    5.36
        Prob > chi2 =    0.0206

( 1)  L._Imotif_3 - L._Imotif_4 = 0

          chi2(  1) =    2.97
        Prob > chi2 =    0.0846

( 1)  L._Imotif_3 - L._Imotif_5 = 0

          chi2(  1) =    0.56
        Prob > chi2 =    0.4534

( 1)  L._Imotif_4 - L._Imotif_5 = 0

          chi2(  1) =    5.87
        Prob > chi2 =    0.0154

( 1)  _IuseXmot_1_2 - _IuseXmot_1_3 = 0

          chi2(  1) =    1.81
        Prob > chi2 =    0.1783

( 1)  _IuseXmot_1_2 - _IuseXmot_1_4 = 0

          chi2(  1) =    0.98
        Prob > chi2 =    0.3228

( 1)  _IuseXmot_1_2 - _IuseXmot_1_5 = 0

          chi2(  1) =   17.81
        Prob > chi2 =    0.0000

( 1)  _IuseXmot_1_3 - _IuseXmot_1_4 = 0

          chi2(  1) =    0.08
        Prob > chi2 =    0.7828

( 1)  _IuseXmot_1_3 - _IuseXmot_1_5 = 0

          chi2(  1) =    2.71
        Prob > chi2 =    0.0996

( 1)  _IuseXmot_1_4 - _IuseXmot_1_5 = 0

          chi2(  1) =    4.57
        Prob > chi2 =    0.0326
```

```
********** wd_avg_log **********

i.used_agent      _Iused_agen_0-1     (naturally coded; _Iused_agen_0 omitted)
i.motif           _Imotif_1-5         (naturally coded; _Imotif_1 omitted)
i.u~ent*i.motif   _IuseXmot_#_#       (coded as above)
Favoring space over speed. To switch, type or click on mata: mata set matafavor speed, perm.
_Iused_agen_1 dropped due to collinearity
_Imotif_2 dropped due to collinearity
_Imotif_3 dropped due to collinearity
_Imotif_4 dropped due to collinearity
_Imotif_5 dropped due to collinearity
Warning: Two-step estimated covariance matrix of moments is singular.
  Using a generalized inverse to calculate optimal weighting matrix for two-step estimation.

Dynamic panel-data estimation, two-step system GMM
------------------------------------------------------------------------------
Group variable: acctid_n                        Number of obs      =    217824
Time variable : fake_year                       Number of groups   =     22606
Number of instruments = 196                     Obs per group: min =         1
Wald chi2(14) =    394.22                                      avg =      9.64
Prob > chi2   =     0.000                                      max =        34
------------------------------------------------------------------------------
             |              Corrected
  wd_avg_log |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
  wd_avg_log |
         L1. |  -.0331673   .0240503    -1.38   0.168     -.080305    .0139705
             |
_Iused_age~1 |  -1.025356   .5192179    -1.97   0.048    -2.043004   -.0077075
   _Imotif_2 |   .3415872   .3683452     0.93   0.354    -.3803562    1.063531
   _Imotif_3 |   .0389254   .4560042     0.09   0.932    -.8548263    .9326771
   _Imotif_4 |    1.12357   .4901612     2.29   0.022     .1628715    2.084268
   _Imotif_5 |   .0532411   .4461203     0.12   0.905    -.8211387    .9276209
             |
   _Imotif_2 |
         L1. |    .166942   .2285219     0.73   0.465    -.2809527    .6148368
             |
   _Imotif_3 |
         L1. |  -.1820273   .2812399    -0.65   0.517    -.7332474    .3691929
             |
   _Imotif_4 |
         L1. |   1.619571   .4324497     3.75   0.000     .7719856    2.467157
             |
   _Imotif_5 |
         L1. |  -.2632907   .3509819    -0.75   0.453    -.9512027    .4246212
             |
_IuseXmot_~2 |   .2532571   .6754374     0.37   0.708    -1.070576     1.57709
_IuseXmot_~3 |   2.350069   .8809798     2.67   0.008     .6233807    4.076758
_IuseXmot_~4 |   .9785519   .9931522     0.99   0.324    -.9679907    2.925095
_IuseXmot_~5 |   1.498035   .7063673     2.12   0.034     .1135809     2.88249
       _cons |   8.080634   .3755708    21.52   0.000     7.344529    8.816739
------------------------------------------------------------------------------
Instruments for orthogonal deviations equation
  GMM-type (missing=0, separate instruments for each period unless collapsed)
    L2.wd_avg_log
Instruments for levels equation
  Standard
    _Iused_agen_1
    _cons
  GMM-type (missing=0, separate instruments for each period unless collapsed)
    DL2.(1b.motif 2.motif 3.motif 4.motif 5.motif)
    DL.wd_avg_log
------------------------------------------------------------------------------
Arellano-Bond test for AR(1) in first differences: z = -17.62  Pr > z =  0.000
Arellano-Bond test for AR(2) in first differences: z =  -1.27  Pr > z =  0.206
------------------------------------------------------------------------------
Sargan test of overid. restrictions: chi2(181)  = 393.62  Prob > chi2 =  0.000
  (Not robust, but not weakened by many instruments.)
Hansen test of overid. restrictions: chi2(181)  = 246.63  Prob > chi2 =  0.001
  (Robust, but weakened by many instruments.)
```

```
( 1)  _Imotif_2 - _Imotif_3 = 0
( 2)  _Imotif_2 - _Imotif_4 = 0
( 3)  _Imotif_2 - _Imotif_5 = 0

          chi2(  3) =    8.83
        Prob > chi2 =    0.0316

( 1)  _Imotif_2 - _Imotif_3 = 0
( 2)  _Imotif_2 - _Imotif_4 = 0
( 3)  _Imotif_2 - _Imotif_5 = 0
( 4)  _Imotif_2 = 0

          chi2(  4) =   10.91
        Prob > chi2 =    0.0276

( 1)  L._Imotif_2 - L._Imotif_3 = 0
( 2)  L._Imotif_2 - L._Imotif_4 = 0
( 3)  L._Imotif_2 - L._Imotif_5 = 0

          chi2(  3) =   29.62
        Prob > chi2 =    0.0000

( 1)  L._Imotif_2 - L._Imotif_3 = 0
( 2)  L._Imotif_2 - L._Imotif_4 = 0
( 3)  L._Imotif_2 - L._Imotif_5 = 0
( 4)  L._Imotif_2 = 0

          chi2(  4) =   30.36
        Prob > chi2 =    0.0000

( 1)  _IuseXmot_1_2 - _IuseXmot_1_3 = 0
( 2)  _IuseXmot_1_2 - _IuseXmot_1_4 = 0
( 3)  _IuseXmot_1_2 - _IuseXmot_1_5 = 0

          chi2(  3) =   10.08
        Prob > chi2 =    0.0179

( 1)  _IuseXmot_1_2 - _IuseXmot_1_3 = 0
( 2)  _IuseXmot_1_2 - _IuseXmot_1_4 = 0
( 3)  _IuseXmot_1_2 - _IuseXmot_1_5 = 0
( 4)  _IuseXmot_1_2 = 0

          chi2(  4) =   13.51
        Prob > chi2 =    0.0090

( 1)  _Imotif_2 - _Imotif_3 = 0

          chi2(  1) =    0.75
        Prob > chi2 =    0.3879

( 1)  _Imotif_2 - _Imotif_4 = 0

          chi2(  1) =    2.90
        Prob > chi2 =    0.0887

( 1)  _Imotif_2 - _Imotif_5 = 0

          chi2(  1) =    1.01
        Prob > chi2 =    0.3140

( 1)  _Imotif_3 - _Imotif_4 = 0

          chi2(  1) =    4.87
        Prob > chi2 =    0.0273

( 1)  _Imotif_3 - _Imotif_5 = 0

          chi2(  1) =    0.00
        Prob > chi2 =    0.9746
```

```
( 1)  _Imotif_4 - _Imotif_5 = 0

          chi2(  1) =    6.03
        Prob > chi2 =    0.0141

( 1)  L._Imotif_2 - L._Imotif_3 = 0

          chi2(  1) =    2.08
        Prob > chi2 =    0.1497

( 1)  L._Imotif_2 - L._Imotif_4 = 0

          chi2(  1) =   22.48
        Prob > chi2 =    0.0000

( 1)  L._Imotif_2 - L._Imotif_5 = 0

          chi2(  1) =    3.40
        Prob > chi2 =    0.0651

( 1)  L._Imotif_3 - L._Imotif_4 = 0

          chi2(  1) =   21.99
        Prob > chi2 =    0.0000

( 1)  L._Imotif_3 - L._Imotif_5 = 0

          chi2(  1) =    0.07
        Prob > chi2 =    0.7946

( 1)  L._Imotif_4 - L._Imotif_5 = 0

          chi2(  1) =   26.50
        Prob > chi2 =    0.0000

( 1)  _IuseXmot_1_2 - _IuseXmot_1_3 = 0

          chi2(  1) =    7.01
        Prob > chi2 =    0.0081

( 1)  _IuseXmot_1_2 - _IuseXmot_1_4 = 0

          chi2(  1) =    0.78
        Prob > chi2 =    0.3766

( 1)  _IuseXmot_1_2 - _IuseXmot_1_5 = 0

          chi2(  1) =    4.90
        Prob > chi2 =    0.0269

( 1)  _IuseXmot_1_3 - _IuseXmot_1_4 = 0

          chi2(  1) =    1.48
        Prob > chi2 =    0.2236

( 1)  _IuseXmot_1_3 - _IuseXmot_1_5 = 0

          chi2(  1) =    0.96
        Prob > chi2 =    0.3279

( 1)  _IuseXmot_1_4 - _IuseXmot_1_5 = 0

          chi2(  1) =    0.35
        Prob > chi2 =    0.5522
```

Bibliography

Holtz-Eakin, Douglas, Whitney Newey, and Harvey S. Rosen. "Estimating Vector Autoregression with Panel Data." *Econometrica* (The Econometric Society) 56, no. 6 (Nov 1988): 1371-1395.

Abakaeva, Julia, and Jasmina Glisovic-Mezieres. *Are Deposits a Stable Source of Funding for Microfinance Institutions?* CGAP Brief, Washington, DC: CGAP, 2009.

Anderson, T. W., and Cheng Hsiao. "Formulation and estimation of dynamic models using panel data." *Journal of Econometrics* 18, no. 1 (Jan 1982): 47-82.

Angrist, Joshua D., and Jörn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist's Companion.* Princeton: Princeton University Press, 2008.

Antonio De Nicolaa, Michele Missikoffa, Roberto Naviglib. "A software engineering approach to ontology building." *Information Systems* 34, no. 2 (2009): 258-275.

Arellano, Manuel, and Olympia Bover. "Another look at the instrumental variable estimation of error-components models." *Journal of Econometrics* 68, no. 1 (1995): 29-51.

Arellano, Manuel, and Stephen Bond. "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations." *The Review of Economic Studies* 58, no. 2 (1991): 277-97.

Arestoff, Florence, and Baptiste Venet. *Learning to Walk Before You Run: Financial Behavior and Mobile Banking in Madagascar*. Document de Travail, Paris: IRD, UMR DIAL, 2013.

Armendáriz de Aghion, Beatriz, and Jonathan Morduch. "Savings and Insurance." Chap. 6 in *The Economics of Microfinance*, by Beatriz Armendáriz de Aghion, & Jonathan Morduch, 147-166. Massachusetts: MIT Press, 2005.

Ashraf, Nava, Dean Karlan, and Wesley Yin. "Deposit Collectors." *Advances in Economic Analysis & Policy* 6, no. 2 (2006).

Ashraf, Nava, Dean Karlan, and Wesley Yin. "Tying Odysseus to the Mast: Evidence From a Commitment Savings Product in the Philippines." *The Quarterly Journal of Economics* 121, no. 2 (2006): 635-672.

Bald, Joachim. *Stability of Small Balance Deposits.* Technical Note, Washington DC: CGAP, 2008.

Ballem, Anjaneyulu, Alphina Jos, Ghiyazuddin A. Mohammad, and Ritika Srivastava. *Making the Business Correspondent (BC) Model Work for Self-Help Groups (SHGs).* MicroSave, 2013.

Baum, Christopher F. "Dynamic Panel Data Estimators." *EC823: Applied Econometrics.* Boston: Boston College, 2013.

Bertrand, Marianne, Sendhil Mullainathan, and Eldar Shafir. "Behavioral Economics and Marketing in Aid of Decision Making Among the Poor." *Journal of Public Policy & Marketing* (American Marketing Association) 25, no. 1 (2006): 8-23.

BFA. *GAFIS Focus Note 2: Stylizing savings behavior to better serve poor clients.* GAFIS Focus Note, Somerville: Bankable Frontier Associates, 2011.

BFA. *GAFIS Focus Note 2: Stylizing savings behavior to better serve poor clients.* GAFIS Focus Note, Boston: Bankable Frontier Associates, 2011.

BFA. *Gafis Focus Note 3: The Impact of Gateway Dynamics on the Business Case for Small Balance Savings.* GAFIS Focus Note, Somerville: Bankable Frontier Associates, 2012.

BFA. *InFocus Note 3: Combining demand and supply side insights to build a better proposition for banks and clients.* InFocus Note #3, Boston: Bankable Frontier Associates, 2012.

Blundell, Richard, and Stephen Bond. "Initial conditions and moment restrictions in dynamic panel data models." *Journal of Econometrics* 87, no. 1 (1998): 115-143.

Bold, Chris, David Porteus, and Sarah Rotman. *Social Cash Transfers and Financial Inclusion: Evidence from Four Countries.* Focus Note 77, Washington, DC: CGAP, 2012.

Browning, Martin, and Annamaria Lusardi. "Household Saving: Micro Theories and Micro Facts." *Journal of Economic Literature* 34, no. 4 (December 1996): 1797-1855.

Buitelaar, Paul, Philipp Cimiano, and Bernardo Magnini. "Ontology Learning from Text: An Overview." In *Ontology Learning from Text: Methods, Evaluation and Applications*, by Paul Buitelaar, Philipp Cimiano, & Bernardo Magnini, 3-12. Amsterdam: IOS Press, 2005.

Burgess, Robin, Rohini Pande, and Grace Wong. "Banking for the Poor: Evidence from India." *Journal of the European Economic Association* (Wiley) 3, no. 2/3 (2005): 268-278.

CGAP. *Developing Deposit Services for the Poor.* Microfinance Consensus Guidelines, Washington, DC: CGAP/The World Bank Group, 2003.

CGAP, AppLab (Grameen Foundation). *Customer Segmentation and Archetypes: Overiew of Processes and Findings in Uganda.* CGAP, Grameen Foundation, July 31, 2012.

Chen, Gregory, and Aimthy Thoumoung. *National Survey of Branchless Banking Agents in India: Towards High Quality BC Customer Service Points.* CGAP, 2012.

Chung, Fu-Lai, Tak-chung Fu, Vincent Ng, and Robert Luk. "Flexible time series pattern matching based on perceptually important points." *International Joint Conference on Artificial Intelligence Workshop on Learning from Temporal and Spatial Data.* 2001. 1-7.

Cohen, Monique. *Building Client Trust and Confidence in Branchless Banking.* Microfinance Opportunities, 2013.

Cohen, Zoe, and Geetha Nagarajan. *Market Segmentation and Outreach to the Poor: Results from Client Welfare Assessment Survey in Peru.* Financial Services Assessment Report, Iris Center, University of Maryland, 2009.

Collins, Daryl, Jonathan Morduch, Stuart Rutherford, and Orlanda Ruthven. *Portfolios of the Poor: How the World's Poor Live on $2 a Day.* Princeton, NJ: Princeton University Press, 2009.

Cracknell, David. *Policy Innovations to Improve Access to Financial Services in Developing Countries: Learning from Case Studies in Kenya.* Center for Global Development, 2012.

Daelemans, Walter, and Marie-Laure Reinberger. "Unsupervised text mining for ontology extraction: an evaluation of statistical measures." *Proceeding of LRE,* 2004: 491-494.

Data and Statistical Services. *Panel Data.* Princeton University. n.d. http://dss.princeton.edu/online_help/stats_packages/stata/panel.htm#stata (accessed 05 12, 2014).

Dean, Debbie. "Savings and the MFI portfolio: Bridging the gap between the needs of the poor and the capability of the institution." *Enterprise Development and Microfinance* (Practical Action Publishing) 22, no. 2 (2011): 118-133.

Deaton, Angus. "Saving in Developing Countries: Theory and Review." *Proceedings of the World Bank Annual Conference of Development Economics 1989.* The World Bank, 1989. 61-96.

Deloitte. *Are we headed towards branchless banking?* Deloitte, 2012.

Demirguc-Kunt, Asli, and Leora Klapper. "Measuring Financial Inclusion: Explaining Variation in Use of Financial Services across and within Countries." *Brookings Papers on Economic Activity* (Brookings Institution Press), 2013: 279-321.

Devedzić, Vladan. "Understanding ontological engineering." *Communications of the ACM* (ACM) 45, no. 4 (2002): 136-144.

Dhawan, Ritesh, Akhand Tiwari, and Vartika Shukla. *Why People Do Not Use Present Banking Systems - A Case for BCs.* MicroSave India Focus Note 63, MicroSave, 2011.

Dolan, Jonathan. *Accelerating the Development of Mobile Money Ecosystems.* Washington, D.C.: IFC and the Harvard Kennedy School, 2009.

Donovan, Kevin P. "Mobile Money, More Freedom? The Impact of M-PESA's Network Power on Development as Freedom." *International Journal of Communications* 6 (2012): 2647-2669.

Dupas, Pascaline, Sarah Green, Anthony Keats, and Jonathan Robinson. *Challenges in Banking the Rural Poor: Evidence from Kenya's Western Province.* NBER Working Paper No. 17851, Cambridge: NBER, 2012.

Eijkman, Frederik, Jake Kendall, and Ignacio Mas. "Bridges to Cash: The Retail End of M-Pesa." *Savings and Development* (Research Center on International Cooperation of the University of Bergamo) 34, no. 2 (2010): 219-252.

Equity Bank Group. *Equity Bank Group third quarter profits soars 26% on the back of growing regional economic activities.* 2015. http://equitybankgroup.com/index.php/investor-relations/financial-results (accessed 02 05, 2015).

Equity Bank Kenya. *Equity Bank Limited Products & Services Tariff Guide.* Nairobi: Equity Bank, 2010.

—. *Equity Ordinary Account.* 2014. http://www.equitybank.co.ke/deposits/equity-ordinary-account (accessed 02 05, 2013).

Equity Bank Kenya. *Investor Briefing: FY 2013 Results.* Annual Report, Nairobi: Equity Bank Group, 2014.

FAI. *Big Questions in Payments.* New York: Financial Access Initiative, 2013.

Ferguson, Michael. *Branchless Banking and Rural Outreach in Malawi: Opportunity International Bank of Malawi's Impact on the Market.* College Park: IRIS Center, University of Maryland, 2011.

FSD Kenya. *FinAccess National Survey 2013: Profiling developments in financial access and usage in Kenya.* Nairobi: Central Bank of Kenya, 2013.

Gilman, Lara, Aneta Genova, and Michelle Kaffenberger. *Mobile Money in the Democratic Republic of Congo: Market Insights on Consumer Needs and Opportunities in Payments and Financial Services.* GSMA, 2013.

Grameen Foundation. *Solutions for the Poorest: Insights on Savings Behavior of Clients.* Grameen Foundation, 2013.

Gray, Bobbi, Jennefer Sebstad, Monique Cohen, and Kathleen Stack. *Can Financial Education Change Behavior?* Working Paper #4, Washignton DC: Microfinance Opportunities, 2009.

Hirschland, Madeline. "Savings Products." Chap. 7 in *Savings Products for the Poor*, by Madeline Hirschland, 137-157. Bloomfield, CT: Kumarian Press, 2005.

Hudon, Marek, and Bernard Liertaer. ""Natural Savings": A New Microsavings Product for Inflationary Environments." *Savings and Development* (Giordano Dell-Amore Foundation) 30, no. 4 (2006): 357-380.

IGD. *Transforming Lives Through Inclusive Financial Access.* IGD Impact Case Study, Washington, DC: Initiative for Global Development, 2013.

Infosys Finacle. *Agency Banking: New Frontiers in Financial Inclusion.* Thought Paper, Infosys Finacle, 2012.

Ivatury, Gautum, and Ignacio Mas. *The Early Experience with Branchless Banking.* Focus Note No. 46, Washington D.C.: CGAP, 2008.

Jack, William, and Tanveet Suri. "Risk Sharing and Transactions Costs: Evidence from Kenya's Mobile Money Revolution." *American Economic Review* 104, no. 1 (2014): 183-223.

Jarrar, Mustafa, and Robert Meersman. "Ontology Engineering - The DOGMA Approach." Chap. 3 in *Advances in Web Semantics I*, by Elizabeth J. Chang, & Katia Sycara. Springer, 2009.

Karlan, Dean, Aishwarya Lakshmi Ratan, and Jonathan Zinman. "Savings By and For the Poor: A Research Review and Agenda." *Review of Income and Wealth* (John Wiley and Sons) 60, no. 1 (2014).

Karlan, Dean, and Jonathan Morduch. *Access to Finance: Ideas and Evidence, The Economics of Saving.* New York, New Haven: The Financial Access Initiative and Innovations for Poverty Action, 2009.

Kaufman, Leonard, and Peter J. Rousseeuw. *Finding Groups in Data: An Introduction to Cluster Analysis.* New York: John Wiley & Sons, Inc., 1990.

Kendall, Jake. *A Penny Saved: How Do Savings Accounts Help the Poor?* Seattle: Bill and Melinda Gates Foundation, 2010.

Lee, Celina. *Agent Banking in Latin America.* Afi Discussion papers, Bangkok: Alliance for Financial Inclusion, 2012.

Levin, Phil, and Gunnar Camner. *Getting the Most Out of Your Data: Segmenting Your Mobile Money Customer Base to Drive Usage.* GSMA, 2013.

Lyman, Timothy R., Gautam Ivatury, and Stefan Staschen. *Use of Agents in Branchless Banking for the Poor: Rewards, Risks and Regulation.* Focus Note No. 38, Washington, D.C.: CGAP, 2006.

Mas, Ignacio, and Hannah Siedek. *Banking through Networks of Retail Agents.* Focus Note. 47, Washington, D.C.: CGAP, 2008.

Mas, Ignacio, and Olga Morawczynski. "Designing Mobile Money Services: Lessons from M-PESA." *Innovations* 4, no. 2 (2009).

Mbiti, Isaac, and David N. Weil. *Mobile Banking: The Impact of M-Pesa in Kenya.* NBER Working Paper Series, Cambridge: National Bureau of Economic Research, 2011.

Mbiti, Isaac, and David N. Weil. "The Home Economics of E-Money: Velocity, Cash Management, and Discount Rates of M-Pesa Users." *American Economic Review: Papers & Proceedings 2013* 103, no. 3 (2013): 369-374.

Midgley, J. "Financial Inclusion, Universal Banking and Post Offices in Britain." *Area* 37, no. 3 (2005): 277-285.

Mileva, Elitza. "Using Arellano-Bond Dynamic Panel GMM Estimators in Stata." Fordham University, July 9, 2007.

Mirkin, Boris. *Clustering for Data Mining: A Data Recovery Approach.* Boca Raton: Chapman & Hall/CRC, 2005.

Morduch, Jonathan. *Household Savings in Developing Economies: An Annotated Reading List.* New York: Financial Access Initiative, April 1, 2008.

N. A., Venkata, and Priyank Mishra. *Success Factors of Equity Bank's Agency Banking.* Microsave Briefing Note # 140, MicroSave, 2013.

Osei-Assibey, Eric. "Financial Exlcusion: What Drives Supply and Demand for Basic Financial Services in Ghana?" *Savings and Development* (Research Center on International Cooperation of the University of Bergamo) 33, no. 3 (2009): 207-238.

Oxford Policy Management. *EFInA: Evaluation of Agent Banking Models in Different Countries.* Oxford Policy Management Ltd., 2011.

Panjwani, Saurabh. "Towards End-to-End Security in Branchless Banking." *Proceedings of the 12th Workshop on Mobile Computing Systems and Applications* (ACM), 2011: 28-33.

Parente, Paulo M. D. C., and J. M. C. Santos Silva. *A Cautionary Note on Tests for Overidentifying Restrictions.* Economics Department Discussion Papers Series, Exeter: University of Exeter, 2011.

Penicaud, Claire, and Arunjay Katakam. *State of the Industry 2013: Movile Financial Services for the Unbanked.* GSMA, 2014.

Pickens, Mark, David Porteus, and Sarah Rotman. *Banking the Poor via G2P Payments.* Focus Note 58, Washington, DC: CGAP and DFID, 2009.

Pickens, Mark, David Porteus, and Sarah Rotman Parker. *Scenarios for Branchless Banking in 2020.* Washington, D.C.: CGAP, 2009.

Platt, Ayn-Byrd, Akhilesh Singh, Sachin Bansal, Anurodh Giri, and Akhand J. Tiwari. *No- Thrills - Dormancy in NFA Accounts.* MicroSave India, 2011.

Radcliffe, Dan. *Expanding Customers' Financial Options Through Mobile Payment Systems: The Case of Kenya.* Global Savings Forum, 2010.

Rangarajan, C. *Report of the Committee on Financial Inclusion.* Mumbai: National Bank for Agriculture and Rural Development, 2008.

Reinberger, Marie-laure, Peter Spyns, Walter Daelemans, and Robert Meersman . "Mining for Lexons: Applying Unsupervised Learning Methods to Create Ontology Bases." *Proceedings of the International Conference on Ontologies, Databases and Applications of Semantics*, 2003.

Reyes, Giovanna Priale, and Denise Dias. *Financial Inclusion and Consumer Protection in Peru.* Joint Assessment Report, CGAP, 2010.

Robinson, Marguerite S. *The Microfinance Revolution: Sustianable Finance for the Poor.* Washington, DC: World Bank, 2001.

Roodman, David. *How to Do xtabond2: An Introduction to "Difference".* Working Paper 103, Washington, D.C.: Center for Global Development, 2006.

Rutherford, Stuart. *The Poor and Their Money.* New Delhi: Oxford University Press, 2000.

Rutherford, Stuart. "Why Do the Poor Need Savings Services?" Chap. 1 in *Savings Services for the Poor*, by Madeline Hirschland, 15-26. Bloomfield, CT: Kumarian Press, 2005.

Schreiner, Mark. "Measuring Savings." *Savings and Development* (Giordano Dell-Amore Foundation) 29, no. 1 (2005): 21-38.

Schreiner, Mark, and Michael Sherraden. *Saving in ADD: Measures from MIS IDA.* CSD Report, St. Louis: Center for Social Development, 2005.

Shah, Shital, Usha Ganesh, and Payal Agarwal. *Discussion Paper I - Identifying Common Operational Challenges facing BCs in the Field.* Discussion Papers On India's Business Correspondent Model, Grameen Foundation USA and Enclude, 2013.

Siedek, Hannah. *Extendng Financial Services with Banking Agents.* CGAP Brief, Washington D.C.: CGAP, 2008.

StataCorp. *Stata 11 Base Reference Manual.* Reference Manual, College Station, TX: Stata Press, 2009.

Tanaka, Yoshiki, Kazuhisa Iwamoto, and Kuniaki Uehara. "Discovery of Time-Series Motif from Multi-Dimensional Data Based on MDL Principle." *Machine Learning* (Kluwer Academic Publishers) 58, no. 2-3 (February 2005): 269-300.

Tiwari, Akhand, Akhilesh Singh, Ann-Byrd Platt, Graham A. N. Wright, Minakshi Ramji, and Sachin Bansal. *The Answer is "Yes" - Cost and Willingness to Pay in India.* MicroSave India, 2011.

Tiwari, Akhand, Ritesh Dhawan, Abhishek Lahiri, Anurodh Giri, Minakshi Ramji, and Sachin Bansal. *Optimising Agent Networks in Gujarat.* MicroSave, 2011.

UCLA: Statistical Consulting Group. *Chapter 2 - Regression Diagnostics.* n.d. http://www.ats.ucla.edu/stat/stata/webbooks/reg/chapter2/statareg2.htm (accessed 11 18, 2014).

Urquizo, Jacqueline. *The Financial Behavior of Rural Residents: Findings from Five Latin American Countries.* Accion, 2012.

V. Banerjee, Abhijit, and Duflo Esther. "The Economic Lives of the Poor." *The Journal of Economic Perspectives* 21, no. 1 (2007): 141-168.

Veniard, Claire. *How Agent Banking Changes the Economics of Small Accounts.* Global Savings Forum, 2010.

Winkler, Gillian M. "Encouraing Savings Behavior through Conditional Cash Transfers: Lessons from Latin America." *Policy Perspectives* 21 (2014): 92-104.

Wise, David A., and Jerry A. Hausman. "Social Experimentation, Truncated Distributions, and Efficient Estimation." *Econometrica* 45, no. 4 (May 1977): 919-938.

Wooldridge, Jeffrey M. *Introductory Econometrics: A Modern Approach.* 5th. Stamford, CT: Cengage Learning, 2013.

World Savings Banks Institute. *WSBI - Key Figures.* 2012. http://www.wsbi.org/template/content.aspx?id=242&LangType=1033 (accessed 09 27, 2012).

Wright, Graham A. N. "Developing Products Based on Market Research." Chap. 6 in *Savings Services for the Poor*, by Madeline Hirschland, 117-135. Bloomfield, CT: Kumarian Press, 2005.

Wright, Graham A.N., and Leonard Mutesasira. *The Relative Risks to the Savings of Poor People.* Briefing Note, Nairobi: MicroSave, 2001.

Zimmerman, Jamie M., and Yves Moury. *Savings-Linked Conditional Cash Transfers.* Policy Brief, Washington, DC: New America Foundation, 2009.

Zimmerman, Jamie M., and Yves Moury. *Savings-Linked Conditional Cash Transfers: A New Policy Approach to Global Poverty Reduction.* A Global Assets Project Policy Brief, New America Foundation, 2009.

Zollman, Julie. *Kenya Financial Diaries: Shilingi Kwa Shilingi - The Financial Lives of the Poor.* Nairobi: Financial Sector Deepening, Kenya (FSDK), 2014.