

A Global Field Experiment on Regulatory Compliance in the Finance Industry

Michael Findley (University of Texas at Austin),
Daniel Nielson (University of Texas at Austin) and
Jason Sharman (University of Cambridge)

NOT FOR FURTHER CIRCULATION

ABSTRACT

This article tests the effectiveness of global banking regulations designed to screen out criminal money based on risk assessment. We did so by setting up shell companies with varying risk profiles to solicit bank accounts directly from banks, and from other intermediary firms. Directly contrary to both the international regulatory regime and scholars' surveyed expectations, both banks and other intermediary firms prove remarkably insensitive to risk in responding for requests for shell company bank accounts. While our strongest risk treatments cause small variations, in the main gate-keepers to the global financial system do not discriminate between high- and low-risk customers. Contrary to baseline rational choice expectations, banks are similarly insensitive to varying rewards. We attribute this insensitivity to the power of organizational scripts: shared, standardized patterns of institutionally constrained agent behavior for diagnosis and response.

This paper provides systematic evidence answering a critical question in international political economy: do banks and firms forming shell companies, the crucial gate-keepers in transnational finance, follow global rules in assessing customer risk before granting access to the international financial system? The risk-based approach is the fundamental regulatory principle in this area. Both regulators and scholars believe the gatekeepers do follow a risk-based approach, but a lack of credible, systematic evidence means that no one really knows. At a time of great stress on the global economic and financial systems, this ignorance is dangerous.

Using a global field experiment, this paper tests the effectiveness of international rules mandating that banks and the firms that set up shell companies adopt a risk-based approach to customers so as to screen out criminals' money from the financial system. It does so by randomly assigning email solicitations to 12,000 banks and other financial intermediaries from every country in the world to one of twelve companies set up by the authors. Because these companies and their solicitations are deliberately created to exhibit sharply varying risk profiles, we can detect respondents' degree of risk sensitivity in the different rates at which they are willing to offer corporate accounts.

Banks are large, highly regulated organizations with complex internal governance hierarchies. Thus, in our pre-registration document and pre-analysis plan we predicted that banks would manage complexity through scripts. Yet these standardized and collective templates are incompatible with the sort of highly differentiated individual judgments the risk-based approach demands of those lower-level employees who decide whether and under what conditions corporate customers can open accounts.

Contrary to the risk-based approach, the central regulatory principle of international banking, and opposed to scholars' expectations as tested in a survey to forecast the results of this study, we find that banks are remarkably insensitive to risk. Working from regulators' manuals, we engineered radically different risk treatments by varying the jurisdiction of company incorporation and the language of the solicitation approach. Yet these varying customer-risk profiles made almost no difference to banks' willingness to offer a corporate account. Banks are also surprisingly insensitive to reward: varying the gross dollar value of business turn-over in the solicitations by two orders of magnitude – from \$500,000 to \$30 million – made no significant difference to outcomes. Firms that provide shell companies are also generally unresponsive to risk.

The significance of this finding is that shell companies with bank accounts are perhaps the single most common mechanism for engaging in money laundering, transnational corruption, tax evasion, and other related crimes (World Bank 2011; FATF 2016, 2018). At a time when the world financial system is under enormous stress, and in which governments have entrusted banks with hundreds of billions of dollars of cheap or free credit, it is more important than ever to find out whether the rules that are meant to regulate banks actually work. Our study raises considerable skepticism on this score.

In substantiating the claims made above, we first provide a brief explanation of the risk at issue: the criminal misuse of bank accounts held through shell companies, the single most important mechanism for tax evasion, money laundering, cross-border corruption, sanction-busting and the like. We then briefly cover the global regulatory regime designed to counter this threat. Regulators have sought to conscript for-profit financial firms into policing the

banking and financial system by imposing a duty to make complicated risk judgements about prospective and actual customers. Shifting from policy to scholarship, the next section explains how as a global field experiment the study advances efforts to improve the external validity of experimental social science. We discuss the subject pool, treatments, outcomes coding, and the expert survey. Finally, we analyze the reasons for and implications of the failure of the risk-based approach and the significance of organizational scripts.

EMPIRICAL AND POLICY BACKGROUND

Because most crime is motivated by profit – from tax evasion to corruption to fraud and human trafficking – criminals face the challenge of moving dirty money through the financial system while remaining undetected. Thanks to globalization, this dirty money increasingly crosses borders. Beyond relatively low thresholds, cash is bulky, conspicuous, and impractical for rapid cross-border transfers, so criminals need to get their money into the banking system. Since the 1980s or 1990s, however, banks have been conscripted into a system of surveillance to prevent illicit wealth entering the financial system. Banks are now mandated to establish the true identity and the riskiness of their clients in connection with a range of crimes (terrorist financing, sanctions-busting, money laundering, tax evasion, etc.). Riskier clients should be subject to greater scrutiny by banks, and if necessary excluded.

For example, although it is not illegal for person to try to open a bank account by visiting a branch with a suitcase of cash, because this behavior is associated with the laundering of the proceeds of crime, it is a high-risk profile. Banks should subject such a prospective customer to extra scrutiny in terms of requiring proof of identity and evidence for the source of wealth, report this behavior to the authorities, and generally impose a higher bar in regulating this individual's access to the banking system. Banks that recklessly pursue profit by taking on risky customers are now routinely hit with multi-billion dollars fines, or in extreme cases deliberately destroyed by sanctions (section 311 of the Patriot Act, Zarate 2013).

A serious obstacle to this goal of banks' knowing their customers has been the use of shell companies, i.e., companies with no substantive business purpose that can be set up online for between a few hundred and a few thousand dollars in a matter of days, but that are nevertheless legal persons that can hold bank accounts and assets. Unless banks know the real person behind the shell company holding the account, the account is *de facto* anonymous, and inward and outward transfers are untraceable. Thus banks must assess the risk of their corporate customers at least as diligently as their individual account-holders.

The global rules in question are set by the Financial Action Task Force (FATF 2012, 2014), the world's anti-money laundering standard-setter and enforcer, an international club comprised of 35 of the world's most powerful countries. The FATF has spawned sister regional organizations promulgating the same rules with memberships incorporating nearly every country. FATF standards in this domain have been endorsed by and incorporated within the standards of additional international organizations, including the United Nations Security Council and International Monetary Fund. They have also been transposed into national legislation. International organizations, national regulators, and private providers produce specialized manuals (e.g. FATF 2019), which we draw from in designing our treatments.

It is important to clarify the relationship between international organizations, states, and firms with regards to these international rules. The standards we are testing have been collectively set down by the member states of the FATF. Individual states are responsible for implementing these rules at the national level, and governments' compliance is periodically and publicly assessed by the FATF. States have legislated that banks in their jurisdiction must follow international rules, centered on the risk-based approach (FATF 2012). In keeping with recent research on Regulatory Intermediary Theory, orchestration, and more general work on global governance, there is thus a chain of different actors involved in applying this international standard (Abbott et al. 2017; Abbott et al. 2015). As noted, if banks fail to comply, they may be sanctioned by national authorities. If too many banks and other firms in a given jurisdiction fail to comply, the national government may be blacklisted and sanctioned by the FATF (Morse 2019).

This regulatory chain linking an international organization to states to firms, in which firms provide "the last mile" of compliance (or non-compliance) with international standards, is common in much other international banking regulation (e.g., Basel III), and a wide range of other international agreements, from aviation safety to environmental standards (Abbott et al. 2017). The bottom line is that if banks and intermediary firms do not assess customer risk, no one else will.

Because governments and international organizations have generally tried to measure the effectiveness of the rules only by reading laws and regulations on the books, rather than by trying to ascertain the standards applied in practice, nobody really knows if they work. Anecdotal evidence, however, gives strong grounds for concern.

Recently, a string of the world's biggest banks have been hit with multi-billion dollar fines, and in some cases criminal convictions, for deliberately or negligently aiding their customers' tax evasion, sanctions busting, money laundering, and rigging of key markets setting inter-bank interest rates and foreign exchange rates. A survey of British banks by the regulator showed that over half failed to perform the mandated extra checks on high-risk customers, with some explaining away lapses because these customers were "from a respectable family" or were "very nice" (FSA 2011: 4, 32). The release of the Panama Papers confirmed that international banks were intimately involved in much of the misconduct revealed (Obermaier and Obermayer 2016). Further leaks in 2020 of banks' private risk reporting to the US Treasury Financial Crimes Enforcement Network (FinCEN) raise worrying questions about the efficacy of the rules designed to keep dirty money out of the financial system. But anecdotes may not aggregate into representative data, so we designed a systematic global experiment to generate unbiased findings in expectation.

THE METHODOLOGICAL CONTEXT: EXPERIMENTS

Proponents of experimental methods have made far-reaching critiques relating to the perils of advancing causal arguments on the basis of observational data (Druckman et al. 2006; Gerber and Green 2012). By conducting a global field experiment, we can test the effectiveness of international rules in a manner that obviates the problems of endogeneity and selection bias that have bedeviled observational research on global governance (Downs et al. 1996; von Stein 2005). In this way, we can get to grips with what in some ways is the fundamental

question of international politics: whether a law-governed order can function in an anarchical international system.

Some scholars have used survey experiments to test public opinion related to international rules (e.g. Tomz 2007). Yet in the context of our study, there are serious problems in using survey experiments to reveal behavior where non-compliance is common, widely seen as inappropriate, or actually illegal (see Findley et al. 2017).

Field experiments are distinguished by the realism or naturalism of the treatments, outcomes, and settings, which together provide confidence in the ability to generalize from the results of the study to the wider world, i.e., stronger ecological validity (Gerber and Green 2012). Because every bank that is connected to the international wire transfer network (SWIFT) is included in our study, the entire population of interest, the study also has uniquely strong external validity. As there is no definitive list of intermediary firms, as they are often unregulated, so it is important to acknowledge that our subjects are a convenience sample of firms with Internet profiles, yet with roughly 7,000 intermediaries from more than 170 countries, the sample has a far wider representation than most experimental studies.

Our field experiment provides other advantages. The random assignment of treatments drawn from regulators' risk-based mitigation manuals enables causal identification and thus well-founded internal validity. Further, participants do not self-select, or even know they are being observed, and thus the study is conducted with a high degree of naturalism in the treatments and outcomes. These features mitigate the danger that participants' responses will be tainted by social desirability bias that could reasonably be expected to arise by simply asking banks or corporate service providers in a survey whether they adhere to due-diligence rules (Findley et al. 2017). As such, together with its coverage of all the world's relevant banks, the naturalism of the exercise, and the fact that participants do not self-select, it has strong external validity compared to most other political science experiments. Thus, the study is well suited to answer big questions about the effectiveness of global rules.

INSTITUTIONAL RESPONSES AND SCRIPTS

Given the claims above, it is vital to establish that our data (correspondence between banks and firms) do in fact represent institutional, rather than individual output. Banks' procedures for taking on new customers are set down in manuals, checklists, training guides, and standard operating procedures precisely to remove individual discretion and set a consistent institutional response. In this sense, though it may be an individual fielding our email (john.doe@hsbc.uk), what we get is an institutional response. Banks spend much time and effort specifically making sure that their policies and procedures are not implemented in an idiosyncratically individual manner; if compliance officers are making decisions on which customers to accept and what identity documents are required on the fly, or according to individual whims, the system has failed. This explains why many email responses were often highly generic, clearly written according to common templates.

Interview material and guidelines produced by international organizations, regulators, and financial institutions specify that the primary unit of interest for corporate accounts is, not surprisingly, the relevant corporation itself. Certainly, establishing the identity of the individual beneficial owner(s) and other corporate officers is crucial in terms of the

international rules, but this fact does not change the inter-institutional character of the relationship. These individuals are of interest precisely because of their role within the institution (company).

Scholars have pointed out the shortcomings in modelling institutions as a single, rational individual writ large, or seeking to explain institutions as simply the sum of individual rational decisions (Powell 2017). Rather than seeking a solution at the level of the individual, according to the psychological maxims of behavioral economics, we change the level of analysis to institutions themselves. Under this view, institutions are more than just the sum of the individuals that make them up, having emergent qualities (March and Olsen 1989).

Furthermore, the logic of decision is much less the deliberate mean-ends calculations of the standard rational expected utility model. Rather than each decision being taken separately, decisions are clustered and categorized into generic classes, and they are linked with equally generic standard responses to form scripts. Partly this process of script-writing is deliberate, and sometimes it may aid efficiency. But just as often it reflects the informal, uncoordinated and anticipated congealing of habits and routines, which may be deeply pathological for organizational function (Barnett and Finnemore 2004). In this sense, institutional scripts are not functional, i.e. they do not exist because of the benefits they may create (Elster 1995).

RESEARCH EXPECTATIONS

The dominant regulatory principle of the risk-based approach bears important similarities with a rationalist account of behavior. Actors should calculate and compare probable utilities associated with alternative courses of action, choosing that which promises the highest expected utility. The alternative courses of action in this case are whether to respond an email solicitation, whether to offer an account, and then, crucially, what level of scrutiny to apply to a prospective corporate customer. The main trade-off is between the riskiness and rewards associated with the customer. As described above, our treatments manipulate risk, and our three levels of business turn-over (\$500,000, \$3 million, \$30 million) manipulate reward, as larger turn-over is associated with higher fees.

To what extent are these incentives institutional, or individual? Both banks and regulators seek to align the two. Individuals bankers may also be prosecuted for excessive risk-acceptance, though it is more likely to lead to being fired. One senior banker described the standard response to such lapses as to “reach out and throttle an underling” (Author interview). Major scandals can lead to whole compliance departments being sacked. Conversely, it is standard practice that bankers who bring in more business or more lucrative business will be rewarded with bonuses and promotion. The net effect of these contrary pressures is essentially unknown – this is the basic uncertainty concerning the effectiveness of regulations that this study seeks to resolve.

In contrast, in terms of institutional scripts, the larger and more regulated the institution, the more likely organizational scripts will shape the behavior of agents. It is commonsensical that while collaboration between just a few individuals can be coordinated on an ad hoc basis, the larger the number of those involved, the greater the diseconomies of scale applying to informal approaches. Governance and coordination of large numbers requires some measure of hierarchy and specialization (Williamson 1975). Socialization in accord with scripts is

much more likely to occur in large, durable organizations than in smaller, ad hoc and transient groupings (March and Olsen 1989).

Given the above, the expectation of a standard rationalist logic and risk-based approach is that banks and bankers will be highly attuned to risk and reward, and hence that the various experimental manipulations will generate large and significant changes in the outcome variables (response rate, refusal, compliance, and non-compliance). In contrast, to the extent that organizational scripts are important, complex judgments weighing risk and rewards will give way to standardized, internally-set responses that are unresponsive to variations in external contacts.

Intermediary firms that set up shell companies and liaise with banks to procure corporate accounts are much smaller, perhaps even just a single person, and less regulated, or in the United States and some other countries, unregulated (Findley et al. 2014). This means that they may have less need for scripts, while they also face different incentives compared to banks. They are much less scrutinized by regulators and are far less likely to attract penalties even when caught out aiding and abetting downright illegal activities (US Senate 2010; Global Witness 2014). A rationalist logic might thus suggest that intermediaries should be relatively insensitive to customer risk, given that they face a much lower likelihood of suffering the consequences of doing business with a corrupt official or money launderer. However in terms of reward, these firms should be at least as sensitive to banks in differential responses to more lucrative prospective customers.

Thus the rationalist, risk-based approach and the organizational scripts accounts have directly opposite predictions on how much banks' and intermediaries' response, refusal, and compliance will vary in response to different risk treatments.

ESTABLISHING EXPECTED RESULTS VIA EXPERT SURVEY

Is it accurate to suggest that the firm baseline presumption for banks is acute sensitivity to customer risk, and hence that findings of script-driven insensitivity to risk would represent a novel and surprising result? Rather than relying on assertion, or a possibly cherry-picked sampling of the literature, we conducted a survey of scholars to establish baseline expectations (see Milgram 1963). We surveyed 116 international political economy scholars.

In the survey, we presented respondents with the study's experimental conditions and asked them to estimate the likely response rates, as well as levels of compliance. More specifically, we asked them to estimate rates across a five-point categorical outcome scale (non-compliance, partial compliance, compliance, refusal, non-response). After this survey and before depositing our pre-analysis plan, we decided to collapse the partial compliance and non-compliance categories (by averaging the responses). However, the main distinction between compliance and partial compliance involved the notarization or certification of the identification documents. We thus softened the definition of compliance to encompass all ID demands of the beneficial owner. For purposes of transparency, we include the original part-compliance and non-compliance responses from the scholars in the survey. See Figure 1 for scholars' expectations of bank behavior across the outcomes. Note that refusals, non-, part-, and full compliance are rescaled as a percentage of anticipated responses. Point predictions

are shown as average predicted proportions; 95-percent confidence intervals are shown by error bars.

As displayed, many of the sizable expected effects are significant statistically compared to the placebo condition. The key takeaway from the scholar survey results accords with the expectation of the rationalist risk-based response model: the strong expectation that banks will be highly sensitive to risk.

Figure 1: Results from Expert Survey Displaying IPE Scholars' Average Predictions of Language Treatment Effects

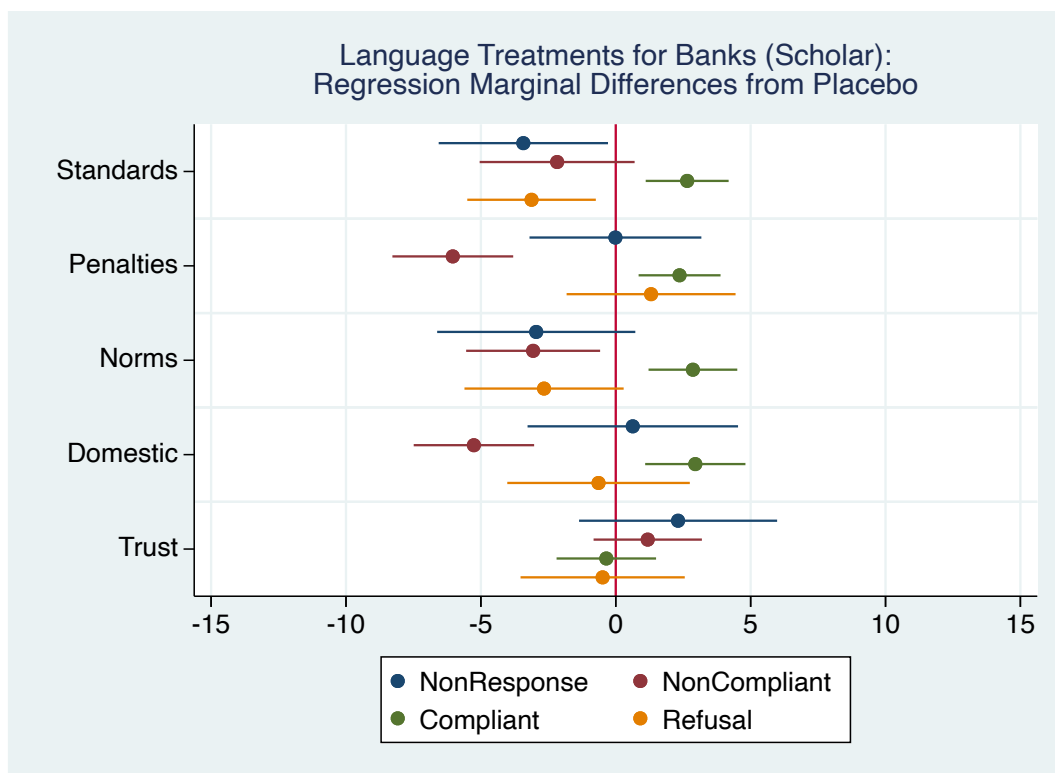
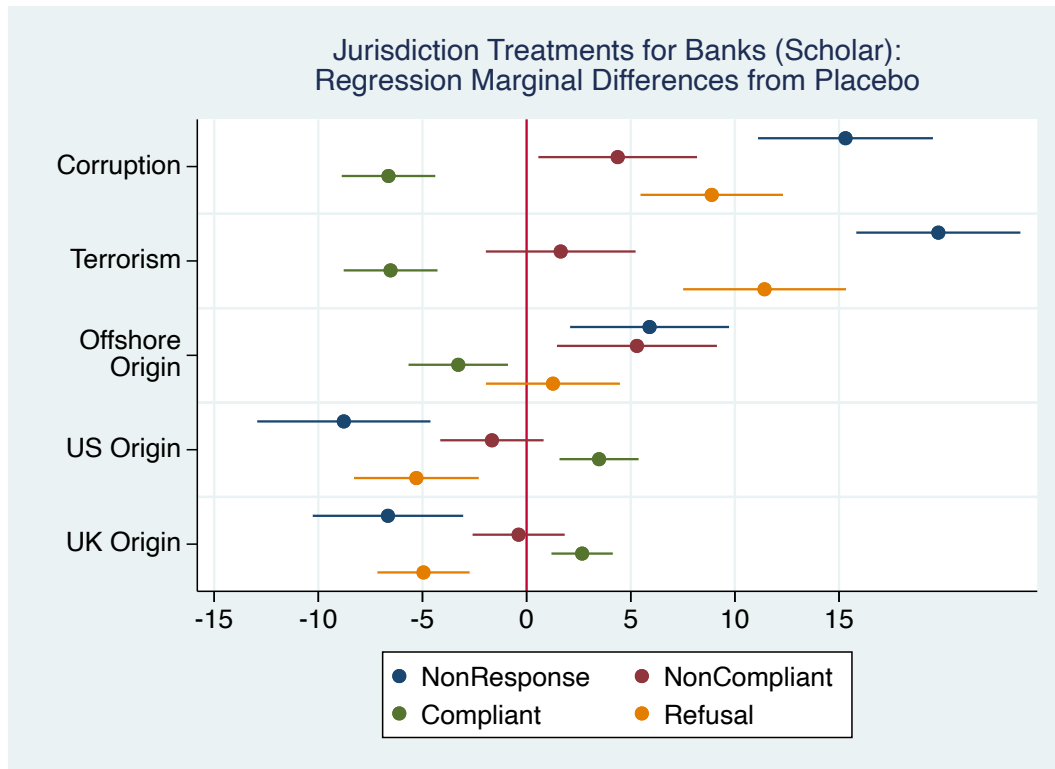


Figure 2: Results from Expert Survey Displaying IPE Scholars' Average Predictions of Jurisdiction Treatment Effects



RESEARCH DESIGN

To probe the extent and variability of scripts across treatment conditions and responding institutions' organizational types, the experiment is based on legally incorporating a variety of companies and having agents working on behalf of these shell companies make email solicitations to thousands of banks and intermediary firms around the world asking to set up a bank account for these companies.

The study does not involve active deception: we have legally incorporated the shell companies seeking bank accounts, and we have a sincere interest in opening bank accounts and making wire transfers. The research design pre-registered the eventual acquisition of multiple bank accounts and the international transfer of funds between them. Many of the banks and intermediary firms will thus add to their profits by working with us and our companies as clients. This sharply diminishes ethical concerns about the potential wasting of banks' and intermediaries' time and effort. Additionally, lying to banks in writing can constitute bank fraud, a criminal offence, so untruths in communication are formally illegal in many jurisdictions. Thus, the experiment was designed to minimize deception, and no false information was involved in the correspondence. It is important to note, however, that researchers did not disclose to banks and intermediaries that the inquiries were part of an academic study. Finally, we delete the names of all individual banks, firms and employees, and otherwise thoroughly de-identify the data, to protect participants from any adverse consequences from the study.

The outcomes of interest are, first, whether banks and other intermediaries reply to our solicitations, second, whether they are willing to offer accounts and, third, whether they follow international rules in verifying the owner of the company. This is in line with our earlier discussion of private firms as the proximate agents of compliance with rules set by international organizations, with states playing an intermediary role linking the two.

The experiment creates different treatments in two ways. The first is to set up shell companies of varying risk profiles, as signaled by the jurisdiction of incorporation. According to international rules, companies and individuals acquire the risk profile of their home countries. For example, a company from a country with a high perceived corruption prevalence should itself be judged a high corruption risk. The second is to insert variations into the language of our approach email to the banks. In expectation, the random assignment of thousands of banks to different company risk profiles and approach emails will enable us to isolate the causes of banks' compliance and non-compliance with international know-your-customer (KYC) rules.

SUBJECT POOL AND NON-RESPONSE CHECKS

The first step in our study was to compile a list of the world's banks and their contact details. The most complete list is that of the Society for Worldwide Inter-bank Financial Telecommunication (SWIFT), the organization that allocates codes for the message system that underpins international bank wire transfers. SWIFT has information on banks (including subsidiaries and branches), in almost every country of the world. Using this list, which comprised our sampling frame, research assistants obtained the relevant email address for roughly 5,000 banks. These 5,000 banks included all headquarter banks on the SWIFT list and a random sample of bank branches in each country. We made three approaches to each bank and branch, for a total of roughly 15,000 approaches, with a wash-out period of three months between each approach to minimize the risk of detection.

The pool of intermediary firms was collected through comprehensive Internet searches. The final sample includes roughly 7,000 firms from more than 200 jurisdictions that set up and sell shell companies and also commonly assist in establishing bank accounts. We likewise made approaches to each of these providers in three randomly segmented tranches of one third each, again with a three-month wash-out period, for a total of approximately 7,000 approaches to the intermediaries.

An important aspect of the findings discussed below is the number of banks and intermediary firms simply not making any response at all to the email solicitations. Ahead of the experiment itself, we sent "non-response" checks to all banks in the sample in order to gauge which banks are responsive at baseline and which are not. It is difficult to sort out exactly what non-response means (a matter we consider in a dedicated section below), the non-response checks allow us to gather some initial information on the banks and also include that information as blocking criterion to better balance the sample. We used five different email scripts, randomly assigned to the banks and intermediaries, for the non-response checks. We asked the firms if they served international customers and if the corporate account could later be upgraded to a merchant account to process credit-card information if needed. The non-response checks revealed that roughly 75 percent of banks and roughly 60 percent of the intermediary firms did not respond to our innocuous inquiry, setting a critical baseline level of subject responsiveness to any email inquiry.

EMAIL APPROACHES

The standard approach template was a message from a representative agent (a researcher) corresponding on behalf of the legally incorporated company, stating the company name and jurisdiction of incorporation. In the approach email, the representative specified that the company is a consultancy concerned with two of five randomly assigned business areas drawn from the actual consulting topics of the beneficial owner, who is one of the principal investigators of the study (the five topics are development assistance, education and training, impact assessment, feasibility studies, and information and communication technology). The RA disclosed the amount of recent business (randomly assigned at “more than” three levels: \$500,000, \$3 million, and \$30 million, the last of which represents the beneficial owner’s total revenue in research grants in the preceding years and thus reflects truthful information). The RA explained that the company would like to establish an account with the bank in the local currency and that the account needed to be able to receive and send international wire transfers. The email then asked how much money establishing the account would cost, how long the process would take, and, crucially for judging rule compliance, what verification documents would be required to set up the account. The email approaches to the intermediary firms requesting bank accounts adopted the same language with slight modifications to match the intermediary context.

It is important to reiterate that all of the information in these base emails was truthful: researchers used their real names; they represented actual, legally incorporated companies based in the stated jurisdictions; and the information presented related truthfully the substance and financing amounts of prior projects by the beneficial owner/author.

Approaches were made via email using specially created email accounts with domain names reflecting the name of the company formed. Company names were generated using random four-letter acronyms and thoroughly vetted to verify that the monikers were not already well known. Communicating via email, rather than phone or video chat, allows for identical treatments, more accurate coding, and a comprehensive record of correspondence.

Email is a standard medium for all but the highest net worth individuals and companies and has been a documented approach across more than 20,000 accounts in one bank alone (US Senate 2014: 83,87,88). We vetted the emails with several parties, including practitioners in the incorporation and banking industries, to ensure that they actually address the key issues of concern to banks and were realistic.

TREATMENTS

The variation between control and treatments arises from the different jurisdictions of incorporation and different language in the approach email. We employ these conditions in both the direct approaches to banks and those via intermediary firms in a fully crossed factorial design involving five information conditions – (1) invoking the international legal standards, (2) noting penalties for non-compliance, (3) referencing norms of compliance, (4) recognizing domestic central bank laws, (5) and demanding banking secrecy – for a combined total of 32 unique conditions. The absence of any of the language conditions is the baseline comparison and was assigned to half of all subject approaches.

We vary the language embedded in the approach email to inform (or, at the very least, prime) subject banks and intermediary firms about international legal standards or a demand for

secrecy. These treatments are designed to test the causal impact of providing information to banks and firms about international law of which they may be unaware or, at a minimum, to test the causal effects of priming the banks and firms about the international standards on which they may have already been briefed. Both of the implications – direct information or prime – are consistent with the argument that information about international standards has a causal effect. Each of the information conditions is assigned independently in a fully crossed factorial design (32 different combinations) with each condition roughly capturing the following ideas.

- Standards: Invokes the requirement that institutions demand documents establishing the identity of the beneficial owner.
- Norms: Expresses interest in following global standards in order to establish and maintain a strong reputation.
- Penalties: Acknowledges that there may be penalties for failure to follow global rules.
- Domestic Enforcement: Notes that the domestic central bank requires identity disclosure and that noncompliance may result in penalties.
- Secrecy: Makes clear that the owner will not disclose his identity because secrecy is of the utmost importance.

We also include several treatments based on the jurisdiction of incorporation. International and national standards mandate that banks assess the riskiness of a potential client in part based on their country of origin. Thus, a company from a country ranked poorly on Transparency International’s Corruption Perceptions Index should be assigned a higher risk than a firm from a low-corruption country, all else equal. Countries have varying risk profiles for money laundering, terrorist financing, being a tax haven, and other concerns. Crucially, according to the international standard for the risk-based approach, high-risk solicitations should get fewer replies, more refusals, and more diligent application of customer due diligence by banks and intermediaries than low-risk solicitations such as our placebo email.

The first treatment is designed to learn whether soliciting offers for bank accounts from companies formed in the United States, the dominant and arguably hegemonic country in the governance of the global economy, affects the response and compliance rate relative to our placebo jurisdictions, Australia and New Zealand. The United States government has been particularly aggressive in applying extra-territorial law enforcement measures and tax regulation to pierce the corporate veil and scrutinize foreign corporate bank accounts. We formed one company each in Delaware and California. Delaware is home to one in three publicly listed US corporations, but it is also often accused of being America’s secrecy haven and leading a race to the bottom in incorporation standards, hence the “Delaware effect.” We formed a second company in California, another of the most popular incorporation jurisdictions in the United States, which has in the past been reputed to have higher standards than Delaware (GAO 2006).

The second set of treatment companies were incorporated in the UK. In addition to London’s being a leading financial center, the British government has sought to take the lead in

enforcing corporate transparency initiatives in the G7 and G20. In 2015, the UK Parliament tightened the country's Know Your Customer regulations.¹

Our offshore treatment aims to learn whether approaches to banks from companies incorporated in stigmatized tax-haven jurisdictions are more or less likely to elicit a compliant response from banks. Offshore centers have been targeted by various multilateral regulatory initiatives and have suffered extensive reputational damage with adverse media coverage. We formed companies in the British Virgin Islands, the most popular single offshore shell company jurisdiction with around 500,000 active companies on its registry, as well as the Seychelles, another stereotypical offshore center hosting more than 100,000 shell companies.

High corruption risk is signaled by companies incorporated in two countries perceived to have major corruption problems: Papua New Guinea (ranked 137th of 180 countries on the 2019 Corruption Perceptions Index) and Bangladesh (ranked 146th). This signal is in accord with the FATF guideline that the risk posed by companies should be in part assessed by the risks in their country of incorporation.

On the same principle that, according to international standards, the country risk is transferred to its corporate citizens, we formed a trust and are still in the process, two years later and counting, of forming a company in Pakistan.² We expected both entities to signal a high terrorism-financing risk. The two conditions also give us leverage on the potential differences in risk between trusts, which are believed to be especially high-risk vehicles, and limited liability companies. Pakistan has earlier played host to the Taliban and leadership of al-Qaeda, as well as a variety of other terrorist groups. Moreover, Pakistan ranks fifth on the 2017 Global Terrorism Index.

The placebo condition was an inquiry originating from either the Australian or New Zealand company and containing only the language from the base email with no additional text from the information treatments. Fully half of all inquiries were assigned to the placebo information condition and one fourth to the placebo jurisdiction condition (another quarter were assigned to the UK companies in anticipation of pooling UK with Australia and New Zealand as an additional placebo jurisdiction). This design choice directly addresses the multiple-comparison problem by increasing the statistical power of comparisons to the control group whose parameters can be estimated with higher precision. The response and compliance rates to the basic placebo control emails form a baseline against which to measure what difference, if any, various risk profiles and information prompts make to banks' willingness to open an account and enforce know-your-customer rules.

Variations in the approach language test the mechanisms and observable implications associated with various theoretical schools of thought (though we hasten to acknowledge that we certainly cannot test the theories themselves directly in their entirety).

¹ We expect that the US and UK treatments will be considered differently by banks and therefore we employ them as separate treatments. With that said, we acknowledge that banks could see these jurisdictions as fundamentally similar in some respects and therefore may end up concluding that we should collapse them for purposes of analysis.

² The fact that the company was still in the process of being formed was disclosed to banks and intermediaries in correspondence.

OUTCOMES AND CODING

Coding was completed on the basis of the email correspondence. The first outcome could be no response at all, which could indicate either disorganization, a commercial judgment that the inquiry is not worth answering, or a form of “soft screening” risk management to turn away undesirable potential clients. Given the high proportion of non-response revealed in the results, we devote a separate section to this in the discussion of the findings.

The second outcome was a refusal to do business, with or without a reason. This was coded as an independent, nominal category.

Third, banks and providers could indicate a willingness to open an account, but require verification documents for the company’s owner. This was coded as compliant. This coding was simplified the documents required by FATF rules (FATF 2012, 2014, 2019). The FATF’s required documents are proof that the company exists (e.g. certificate of incorporation, which shows jurisdiction of incorporation), a copy of the company’s by-laws (articles or memoranda of incorporation), and its registered address. Most important, and therefore key to our coding rules, is proof of the beneficial owner’s identity, evidenced by a verified copy of a government photo identity document, usually a passport, or an in-person visit to the bank to establish identity. If the bank or intermediary required a passport copy but did not specify that it must be for the beneficial owner (or might instead apply to a nominee director or agent, both of which can enable anonymous banking), researchers followed up to reduce the ambiguity. Banks and firms coded as non-compliant failed to specify required ID of the beneficial owner after repeated follow-up.

Banks and providers that are willing to open an account without this supporting documentation material to establish the owner’s true identity were thus coded as non-compliant, as they are breaking international rules by offering what amounts to an anonymous or untraceable bank account. Without ID of the beneficial owner, it is difficult or impossible to find the real person who controls the account via the company and thus that person can commit financial crimes with impunity.

Hence there are four outcomes: no response, refusal, compliant, and non-compliant. Emails were independently coded by two research assistants for accuracy and consistency, with discrepancies adjudicated by a third, senior coder.

DATA ANALYSIS

As pre-registered and deposited prior to data collection in our pre-analysis plan, initial data analysis employs difference-in-proportions tests comparing each treatment condition to the control. Differences can be assessed with t-tests assuming the t distribution. Additionally, we employ multinomial probit analysis to consider treatment effects alongside covariates on all four outcomes simultaneously.

We also report selection models. The logic behind the selection model is that three of the outcomes – compliance, non-compliance, and refusal – can only be observed if a subject responds. As such, response acts as a hurdle or gate-keeping stage in estimation. In our case,

we do not have different variables to identify separate stages of a selection model, so we draw on the estimation technique developed by Sartori (2003), which allows the same variable to identify both stages.

Below we discuss the three most important elements of our findings: the high proportion of non-responses, the uneven effectiveness of the Know Your Customer rule, but most especially the general (although not complete) insensitivity to risk.

NON-RESPONSE: SOFT-SCREENERS OR NEVER-RESPONDERS?

Excluding auto replies, the response rate for banks to our email solicitations was 22.7 percent (i.e. including compliant, non-compliant and refusal responses), 3,532 responses from 15,563 solicitations. The earlier response checks had predicted roughly this level of replies. Nevertheless, the high rates of non-response demand an explanation, both on their own account, and in terms of what it means for the conclusions that can be drawn from the replies that we did get.

The non-response category poses significant inferential challenges because non-response could be capturing a wide array of information. The most important question is whether the banks and firms not replying to the solicitations were “soft-screeners,” i.e. deliberately managing risk by taking decisions to ignore the messages, or “never-responders,” who simply never respond to any international email approaches, perhaps because they are unable or unwilling to cater to foreign customers. Soft-screeners are relevant to our study because they could potentially provide access to the international banking system in a compliant or non-compliant manner. Even if they did not reply to our solicitations, other different approaches from prospective foreign clients might induce them to do so. On the other hand, because “never responders” never provide access to the international banking system, by definition they are not relevant to a test of the conditions under which it is possible to access to that system. While a full understanding of non-response is elusive, we did take some additional steps which suggest that the large majority of the non-responses are “never-responders,” and hence not a relevant population for the aims of this study. If so, this ameliorates, but does not entirely solve, the challenges to interpreting results.

As noted, ahead of the experiment itself, we sent non-response checks to all banks in the sample in order to gauge which banks are responsive at baseline. The non-response checks allow us to gather some initial information on the banks and use that information as blocking criteria to better balance the sample. We used five different email scripts, randomly assigned to the banks, for the non-response checks. We asked the banks if they served international customers and if the corporate account could later be upgraded to a merchant account to process credit-card information if needed. The non-response checks revealed that roughly 75 percent of banks and 55 percent of intermediaries did not respond to our innocuous inquiry.

Also prior to the launch of the experiment, we randomly sampled 95 of the banks that did not respond to these non-response checks and called them by phone to seek to understand why they failed to reply. Many banks lacked a phone number and we were unable to contact them. For the rest, we used a Google Voice number to contact each bank, and each were called between Monday and Friday, from 9 am to 5 pm in the institution’s local time. When contacted, many calls continued to ring with no answer. Several banks had the call dropped

before any contact was made. A large majority of banks (over 75 percent) could not be contacted on the phone number they provided on their website.

Of the banks that were contacted successfully, several had logistical issues, such as a language barrier when English was not the primary language used among the staff. We received a few responses where the bank chains claimed inability to open accounts internationally or required in-person verification. They were insistent that these were new international restrictions placed on the institution. These banks were located in Papua New Guinea, Switzerland, and Austria. The phone call approach primarily revealed that many international banks did not answer their phones.

SCRIPTS NOT RISK: THE FAILURE OF THE RISK-BASED APPROACH

Beyond the baseline result of having a viable number of replies, the second major endorsement of the research design is that in some instances banks did react to at least some treatments, though significant effects for information treatments (see Figure 2) are few and are marginally weaker than for jurisdiction treatments (see Figure 3). In this sense, enough banks are reading and responding to the different emails to show some treatment effects.

Figure 3: Effects on Bank Responses from Language Treatments (asking Banks for Bank Accounts)

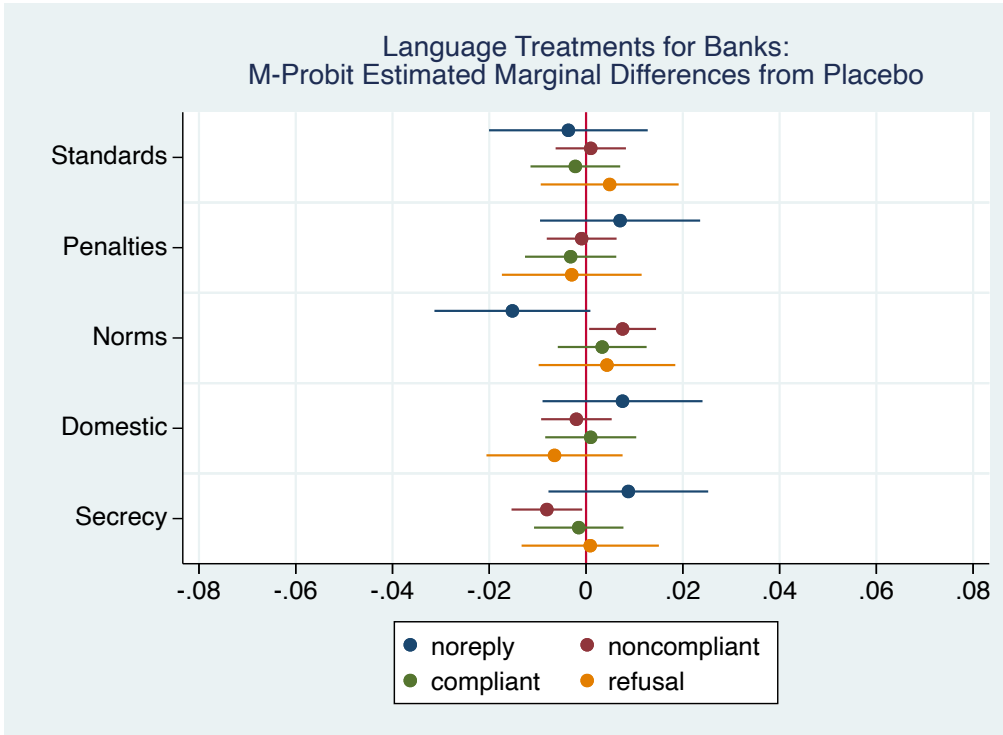


Figure 4: Effects on Bank Responses from Jurisdiction Treatments (asking Banks for Bank Accounts)




Figure 5: Effects on CSP Intermediaries' Responses from Language Treatments When Asking for Bank Account Only




Figure 6: Effects on CSP Intermediaries' Responses from Jurisdiction Treatments When Asking for Bank Account Only

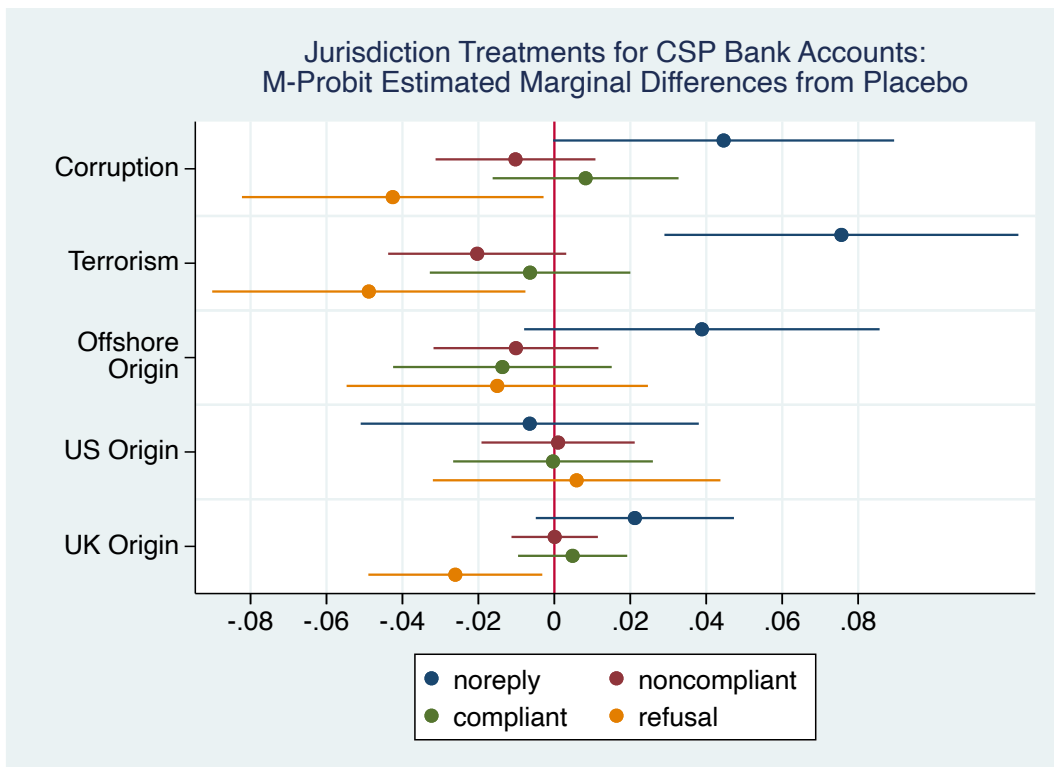


Figure 7: Effects on CSP Intermediaries' Responses from Language Treatments When Asking for both Bank Account and CSP (Hybrid)





Figure 8: Effects on CSP Intermediaries' Responses from Jurisdiction Treatments When Asking for both Bank Account and CSP (Hybrid)



Nevertheless, with only minor exceptions covered below, the major conclusion is that banks are surprisingly indifferent to customer risk. With only rare and generally minor exceptions, even strong and obvious differences in customer risk profile generally did not make a significant difference to banks' response, refusal, and compliance rates. Even where there was a significant difference, the substantive effect size was very small, usually in the range of 1-3 percent. We know that this finding is indeed surprising to scholars, because our expert survey, as reported in Figures 1 and 2, clearly predicted both many more significant treatment effects, and much larger substantive treatment effects than were actually observed, especially for jurisdiction treatments.

Though we did not survey regulators, the experimental results cut directly against the central principle of the global regulatory regime designed to combat transnational financial crime: as a rule, banks and intermediary Corporate Service Provider firms do not assess customer risk in granting or withholding shell companies access to the international banking systems.

This relative insensitivity to risk is in line with the expectations of institutional scripts: the idea of generic, standardized responses to fulfilling recurrent and repetitive organizational tasks, such as the procedure for opening new corporate accounts, that are resistant to the imposition of finely-grained differentiation and considerations of individual judgment. Regulators have sought to transfer complex and delicate questions of differentiating between customers to establish different levels of customer risk and differential treatment to banks. The evidence in favor of scripts suggests that regulators have mostly failed in this objective.

For information treatments, most are associated with higher average Non-Response, with Secrecy however the only condition that caused a statistically significant increase in non-response (and then only at the 0.1 level). For the remaining outcomes, Secrecy is associated with significantly higher Refusal but also significantly lower Compliance and Non-Compliance. Secrecy appears to be the only information condition that causes significant outcome differences in bank responses, yet it is important to reiterate that even here the substantive effect size is very small.

Similarly, the highest risk jurisdiction treatment, Terrorism, caused statistically significant higher Non-Response, an increase in Refusals, and decreases in Compliance and Non-Compliance. As with Secrecy, we interpret this as banks that shifted from the latter two categories toward refusal of the high-risk potential client. All of the jurisdiction conditions are associated with significantly higher Non-Response compared to the placebo condition of Australia/New Zealand. Yet only the Offshore treatment of British Virgin Islands or Seychelles caused a significant difference in the other outcomes and there only in decreasing the rate of Compliance – not necessarily an auspicious effect for what many would see as higher-risk customers.

Thus, both high-risk treatments – Secrecy and Terrorism – elicit parallel responses from banks: significantly increasing Non-Response and Refusals while decreasing Compliance and Non-Compliance, but with the substantive effect sizes being small. This suggests that both treatments shift banks toward higher rejection and away from both permissive and restrictive replies. Together with the fact that the treatments did not make a difference to the level of Compliance, this suggests that to the extent there is any sensitivity in the system, banks respond to risk by refusing would-be high-risk customers, and they make a simple binary choice to accept or reject potential clients, rather than adjusting their propensity to conduct

Know Your Customer checks. This is in contrast to the reaction that might be expected, whereby high-risk customers would get more scrutiny relative to low-risk customers. Thus it appears that banks do not apply a risk-based approach to customer due diligence, KYC procedures roughly constant despite widely varying levels of customer risk. This finding is especially noteworthy remembering that our treatments are generally directly derived from the FATF checklist of risk indicators that are meant to set the standards for banks; to this extent, they are textbook, obvious risks and hence strong treatments.

Aside from our language and jurisdiction treatments, we also randomized the amounts of money our shell companies specified for their turn-over. The three figures provided: \$500,000, \$3 million or \$30 million over five years (the authors have in fact been involved with projects of these values, in accord with our no-deception rule for dealing with banks). The intuition here is fairly straight-forward: banks may treat rich customers differently than poorer ones. The greater attractiveness of a \$30 million turn-over firm might mean that banks are keener to accept such clients (lower Refusal and lower Non-Response). They might raise Compliance (as it is more worthwhile for banks to go to the trouble to conduct due diligence checks), or raise Non-Compliance (banks are more inclined to accept more risks in flouting the rules for greater rewards).

In fact, however, we found that even across these three orders of magnitude the potential value of the customer made no significant difference to the results. This finding once again suggests that banks are comparatively unresponsive to risk and reward in deciding whether or not to accept customers, and whether or not to apply Know Your Customer standards.

Banks' insensitivity to reward is important in helping disconfirm another potential explanation: perhaps banks have no inherent interest in screening out criminals (after all they are profit-makers, not law enforcers), and are willing to trust their luck in surreptitiously defying regulators? Such an explanation might just accommodate accepting likely criminals at the same rate as low-risk corporate customers, but it is much harder to square with banks' tendency to accept high-rollers at the same rate as firms with a much small turn-over.

While it appears that banks proved relatively insensitive to risk, the results suggest that corporate service provider (CSP) intermediaries prove even less sensitive, as seen in Figures 5-8. This applies both to requests for a bank account for one of our existing companies (Figures 5 and 6), and Hybrid requests for both a shell company and a bank account to go with it (Figures 7 and 8). While the Secrecy treatment caused significant increases in Non-Response and (marginally) in Refusals and caused a concomitant decrease in Non-Compliance, it left the Compliance rate relatively unchanged. Among the other information conditions compared to placebo, only the Standards condition caused a meaningful difference in a single outcome measure by significantly decreasing the Non-Compliance rate.

The jurisdiction conditions also produced relative insensitivity among intermediaries. The highest-risk Terrorism treatment caused a small increase in Non-Response but left the other outcomes relatively unchanged in any significant way. The Corruption treatment caused a significant decrease in the Non-Compliance rate but did not appreciably alter the other outcomes. Thus, contrary to researcher expectations, it appears that intermediaries, like banks, are relatively insensitive to customer risk.

CONCLUDING THOUGHTS ON COMPLIANCE

A final matter concerns the baseline question of what counts as compliance. An important challenge for both our coding decisions and the contemporary policy debate inside the Financial Action Task Force concerns the “who” and “how” of establishing the identity of shell company owners. As discussed, the critical question for ensuring the transparency of shell companies and the corporate accounts held in their name is identifying the beneficial (i.e. real) owner. The most obvious implication from this is the need to verify that the customer is who she or he says they are with reference to government documents, typically a passport. Yet for the transparency rules to work, banks and other intermediaries must not only verify a person’s identity, but also verify that that person is in fact the beneficial owner, rather merely than a stand-in or proxy for the real owner. Thus as well verifying an individual, for the system to work there is also a requirement to verify that individual’s relationship to the company, i.e. to verify that they are indeed the beneficial owner.

The significance of this double requirement, not just the “who” but also the “how” of beneficial ownership, is significant for our research design, but also currently a central question for policy-makers. Our solicitation emails are sent out on behalf of companies by research assistants who represent, but do not own, the companies. A reply from a bank or firm that “We need to verify your identity with a passport copy” raises the question of *whose* identity needs to be verified: the representative, which counts as a non-compliant response, or the owner, which is a compliant response?

For policy-makers, getting both parts of beneficial ownership verification right is no mere quibble, but rather goes to the heart of the effectiveness of the regime. In setting up shell companies, it is standard that key corporate roles like company director or secretary will be filled by a nominee, a professional stand-in. The nominee individual may be acting in this role for hundreds of companies without having any idea of what those companies actually do or even who owns them. Indeed, this nominee or corporate proxy function is one of the main services offered by intermediary firms that set up and service shell companies, and connect them with corporate bank accounts. As such, even the most thorough and searching identity checks will be for naught if they applied to the wrong person, i.e. the nominee rather than the actual owner.

In our study, about of third of the responses which specified the need for identity documents did not specify that these documents had to relate to the actual owner. Either explicitly or implicitly, they allowed that verification of the representative would be sufficient. Aside from being a worrying policy vulnerability, this may also feed back into our scripts story. Obtaining a passport copy for every applicant for an account is a relatively easy requirement to standardize, being a box to be checked. But ensuring that this person is the actual beneficial owner is more difficult, given that there are many different ways this ownership can be exercised.

BIBLIOGRAPHY

- Abbott, Kenneth W., Philip Genschel, Duncan Snidal and Bernhard Zangl (Editors) (2015) *International Organizations as Orchestrators*. Cambridge: Cambridge University Press.
- Abbott, Kenneth W., David Levi-Faur and Duncan Snidal (2017) “Introducing Regulatory Intermediaries.” *Annals of the American Academy of the Political and Social Sciences* 670 (March): 6-13.
- Basel Committee on Banking Supervision. 2012. *Core Principles for Effective Banking Supervision*. Basel, Switzerland: Bank for International Settlements.
- Downs, George W., David M. Roake and Peter N. Barsboom (1996). “Is the Good News about Compliance Good News about Co-operation?” *International Organization* 50 (3): 379-406.
- Druckman, James N., Donald P. Green, James H. Kuklinski and Arthur Lupia (2006). “The Growth and Development of Experimental Methods in Political Science.” *American Political Science Review* 100 (4): 627-635.
- Financial Action Task Force (FATF) (2012). *International Standards on Combating Money Laundering and the Financing of Terrorism & Proliferation* (Paris).
- Financial Action Task Force (FATF) (2014). *Transparency and Beneficial Ownership* (Paris).
- Financial Action Task Force (FATF) (2016). *FATF Report to the G20: Beneficial Ownership* (Paris).
- Financial Action Task Force (FATF) (2018) *Concealment of Beneficial Ownership* (Paris).
- Financial Action Task Force (FATF) (2019) *Guidance for a Risk-Based Approach for Trust and Corporate Service Providers (TCSPs)* (Paris).
- Findley, Michael G., Daniel L. Nielson, and J.C. Sharman (2014). *Global Shell Games: Experiments in Transnational Relations, Crime and Terrorism* (Cambridge: Cambridge University Press).
- Findley, Michael G., Brock Laney, Daniel Nielson, and J.C. Sharman (2017). “External Validity in Parallel Global Field and Survey Experiments on Anonymous Incorporation.” *Journal of Politics* 79 (4): 856-872.
- Financial Services Authority (FSA). 2011. *Banks’ Management of High Money Laundering Risk Situations*. London.
- Gerber, Alan S. and Donald P. Green (2012). *Field Experiments: Design, Analysis and Interpretation* (New York: W.W. Norton).
- Global Witness. (2014). *The Great Rip-Off: Anonymous Company Owners and the Threat to American Interests*. Washington D.C.
- Government Accountability Office (2006). “Company Formations: Minimal Ownership Information is Collected and Available” (Washington).
- March, James G. and Johan P. Olsen (1989). *Re-discovery Institutions: The Organizational Basis of Politics*. New York: Free Press.

- Milgram, Stanley. 1963. "Behavioral Study of Obedience." *Journal of Abnormal and Social Psychology* 67(4): 371-378.
- Morse, Julia. 2019, "Blacklists, Market Enforcement, and the Global Regime to Combat Terrorist Financing." *International Organization* 73 (3): 511-545.
- Obermayer, Bastian and Frederik Obermaier (2016). *The Panama Papers: Breaking the Story of How the Rich and Powerful Hide their Money* (New York: Oneworld).
- Powell, Robert. 2017. "Research Bets and Behavioral IR." *International Organization* 71(S1): S265-S277.
- Tomz, Michael. (2007). "Domestic Audience Costs in International Relations: An Experimental Approach." *International Organization* 61 (4): 821-40.
- US Senate Permanent Subcommittee on Investigations. 2014. "Offshore Tax Evasion: The Effort to Collect Unpaid Taxes on Billions in Hidden Offshore Accounts."
- von Stein, Jana (2005). "Do Treaties Screen or Constrain? Selection Bias and Treaty Compliance." *American Political Science Review* 99 (4): 611-622.
- Williamson, Oliver E. (1975) *Market and Hierarchies: Analysis and Antitrust Implications*. New York: Macmillan.
- World Bank (2011). *The Puppet Masters: How the Corrupt Use Legal Structures to Hide their Stolen Assets and What to do About It* (Washington D.C.: World Bank).
- Zarate, Juan C. (2013) *Treasury's War: The Unleashing of a New Era of Financial Warfare* (New York: Public Affairs).