Public versus Private Protection against Crime

The Case of Somali Piracy

Quy-Toan Do

Lin Ma

Claudia Ruiz

WORLD BANK GROUP

Development Research Group

Poverty and Inequality Team & Finance and Private Sector Development Team

July 2016

Abstract

This paper analyzes the rise and demise of piracy off the coast of Somalia to highlight the interplay between public and private crime protection measures. Using unique data on attacks, hijacks, and ransoms, the authors estimate a structural model of Somali piracy to calculate the elasticity of crime with respect to two forms of protection: the publicly-provided deployment of international navies

in the Western Indian Ocean and the private provision of maritime security contractors on board vessels. They assess both positive and negative spillovers associated with private security. Their findings have implications for the optimal composition of protective measures and the regulation of the market for private security..

This paper is a product of the Poverty and Inequality Team and the Finance and Private Sector Development Team, Development Research Group. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at <http://econ.worldbank.org>. The authors may be contacted at abonfield@worldbank.org.

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

Public versus Private Protection against Crime: The Case of Somali Piracy*

Quy-Toan Do

Lin Ma

Claudia Ruiz

World Bank

National University of Singapore

World Bank

Keywords: Crime, protection, deterrence, displacement, externalities

JEL Codes: K42, O12

*We are grateful to Abhijit Banerjee, Jishnu Das, Sebastian Galiani, Clément Gorrissen, Hanan Jacoby, Clément Joubert, Ted Leggett, Andrei Levchenko, Dilip Mookherjee, Ben Olken, Martin Ravallion, Debraj Ray, and seminar participants at the 31st BREAD Conference (Georgetown), Duke, Georgetown, Harvard and MIT, Paris School of Economics, University of Maryland, and the World Bank for helpful discussions. Special thanks go to Cyrus Mody and the IMB for providing us with the dataset on piracy incidents. The findings, interpretations, and conclusions expressed in this work do not necessarily reflect the views of the World Bank, its Board of Executive Directors, or the governments they represent.

1 Introduction

The dramatic rise in pirate attacks off the coast of Somalia over the years 2005-2011 stunned the international community. In total, 1,101 reported attacks have been attributed to Somali pirates, resulting in the hijacking of 217 vessels and their crews and an estimated US\$338m paid in ransoms. Equally astonishing was its collapse in 2012, which has been largely sustained until now. The eradication of Somali piracy has been attributed to the deployment of naval assets in the Western Indian Ocean, celebrated as a rare show of international military collaboration, and to the increased use by shipping companies of private security contractors on board vessels sailing through pirate-infested waters.

In this paper, we focus on the tension between the two aforementioned forms of protection against crime, one publicly- and the other privately-provided, which led to its demise. To do so, we set up a structural model of Somali piracy, which we estimate using novel and rich data on incidents attributed to Somali pirates. On the one hand, the effect of international navy patrols is found to have been modest overall. On the other hand, the private provision of security was instrumental and characterized by both deterrence and crime displacement externalities.

The Somali piracy business model that we look at is exclusively a kidnap-and-ransom model, whereby the sole purpose of hijacking a vessel is to “sell” her back in her entirety to the shipping company as opposed to cargo and vessel theft, more prevalent elsewhere in the world. Attacks are launched against vessels sailing through the Gulf of Aden and beyond. When successful, vessel and crew are brought back to Somalia’s shores and a separate team takes over to protect and maintain the ship and feed the crew while negotiation with the shipping company for their liberation takes place.

Therefore, we model pirates as infinitely-lived agents who, in each period, decide whether to launch an attack against ships sailing off the coast of Somalia and when they do, choose what vessel to target. If successful, a bargaining game with asymmetric information is played between the pirates and the shipping company, which determines duration of captivity and ransom amount. Protective measures against crime come in two forms: navy patrols are a public good in that they uniformly lower the probability

that an attack is successful, while private maritime security contractors only protect the vessel on which they are deployed.

Our model delivers a number of time-varying equilibrium outcomes that are functions of structural parameters of the model: the number of pirate attacks and the success rates, the average size of the vessels attacked and hijacked, and for the latter, the ransom extracted and the duration of captivity. These structural parameters are estimated using the method of simulated moments. We use a unique data set that contains, for each incident attributed to Somali pirates over the 2005-2012 period, its date, the characteristics of the vessel that was attacked, whether the attack resulted in the hijacking of ship and crew, and if so, the duration of captivity and amount of ransom paid.

We can then construct counterfactual simulations to measure the elasticity of crime with respect to the level of protection. In doing so, we quantify the mechanisms behind the effects that the different types of security interventions have on the evolution of Somali piracy. We find that navy patrols reduce the success rates of pirates across all types of ships, therefore deterring entry into the business in the first place. Their overall effect is however found to be small. On the other hand, the deployment of private armed security contractors on board vessels explains most of the fall in pirate attacks. The deployment of armed guards among larger vessels lowers their hijacking success probabilities and is found to have two effects. First, as with the navies, private maritime security contractors have a deterrence effect. Such positive externality benefits *all* vessels. Specifically, we find that if an additional one percent of the largest vessels sailing off the coast of Somalia in 2011 had been equipped with armed security personnel, it would have resulted in a 35.3 percent drop in overall attacks. This positive spillover is similar to the one identified in Ayres and Levitt (1998). Second, as commonly highlighted in the crime literature, protection also induces a displacement of crime towards more vulnerable targets, the smaller ships in this context. Despite the aforementioned deterrence externality, we estimate that the same one percent increase in the incidence of armed security on board largest vessels leads, through crime displacement, to a 22.1 percent *increase* in attacks against smaller vessels.

Finally, we simulate various scenarios that look at the elasticity of crime with respect

to both the level and composition of protective measures. In particular, our policy simulations vary the composition of security interventions to determine the marginal rate of substitution between navy patrols and private maritime security contractors. Using cost estimates, we discuss the optimal allocation of limited funds between these two crime protection instruments. Given the deterrence and displacement externalities generated by private security, we contrast the optimal protective regimes obtained when the objective is to minimize overall crime incidence only versus when the distributional impacts of private protection are considered. In doing so, we emphasize the role of navy assets deployed on the Indian Ocean in addressing the implications of crime displacement.

Our paper directly relates to the few economic analyses of piracy, and notably to Ambrus, Chaney and Salitskiy (2018). The authors analyze data from Renaissance-era piracy activity in the Mediterranean sea to estimate a model of delay in bargaining. While our paper adopts a similar approach to link ransoms and durations of captivity, we also model other decisions such as entry and vessel choice. The main objective of our paper is to capture enough features of the piracy business to be able to discuss the role and interaction of security measures.

Our contribution, therefore, is an addition to the literature on the economics of crime pioneered by Becker (1968) and Ehrlich (1973), and recently reviewed by Draca and Machin (2015). One outcome of our analysis is the estimation of the crime-policing elasticity, where policing here is achieved by a mix of navy patrols and private maritime security contractors. We thus complement the findings of Levitt (1997), Corman and Mocan (2000), Di Tella and Schargrodsky (2004), or Draca, Machin and Witt (2011). A noteworthy difference is that the identification of the impact of protection on crime in these and most other studies relies on arguably exogenous spatial heterogeneity in the level of policing. Such quasi-experimental settings generate a “natural” counterfactual, which allows identifying crime-policing elasticities. These studies, however, focus on a single protection instrument and thus cannot speak to the central question of this paper, which is the tension between public and private policing. In our setting, there is no such natural counterfactual, since all pirates are subject to the same policing environment. We instead estimate a structural model of the piracy business. By putting structure on agents’ behavior, our model

allows not only constructing a counterfactual to measure the elasticity of crime with respect to protection, but also evaluating the effects and interactions of multiple protection instruments. In particular, we are able to gauge the externalities generated by the private provision of security (Clotfelter 1978, Cornish and Clark 1987, Di Tella, Galiani and Schargrodsky 2010, Vollaard and van Ours 2011, van Ours and Vollaard 2015, Ayres and Levitt 1998, Gonzalez-Navarro 2013, Banerjee, Duflo, Keniston and Singh 2012), which in most natural experiment settings are a potential source of concern for identification.

The central policy discussion in the paper pertains to the optimal allocation of police resources as well as the nexus between public and private provision of protection, which was discussed by Clotfelter (1977) and has since benefited from the more recent contributions by e.g. Kremer and Willis (2016) and Galiani, Cruz and Torrens (2018). Finally, the discussion of publicly-provided navy patrols versus private maritime security contractors can also be viewed as the analysis of targeted versus un-targeted police interventions (Lazear 2004, Eeckhout, Persico and Todd 2010, Banerjee et al. 2012). However, the targeted interventions in our setting are characterized by deterrence externalities, making them essentially un-targeted.

The rest of the paper is organized as follows. Section 2 provides background information on Somali piracy and presents some stylized facts that will guide our modeling choices. In section 3, we lay out a model of Somali piracy, which we estimate in section 4. We conduct our policy simulations in section 5. Section 6 concludes.

2 Piracy off the Coast of Somalia

While some attacks by pirates from Somalia have been reported earlier, the onset of piracy as kidnapping-and-ransom can be dated to the hijacking of the MV Feisty Gas in April 2005. That year, a total of 45 pirate attacks were reported to the International Maritime Bureau (International Maritime Bureau 2006). In 2006, the number of attacks dipped compared to the year before, most likely due to an attempt by the Islamic Courts Union (ICU), an Islamist administration competing for power with the Transitional Federal Government (TFG), to eradicate piracy it had deemed to be anti-Islamic (World Bank 2013). The

fall of the ICU in 2006 was followed by a period of relative stability as a weak TFG left the central and northern regions of Somalia under control of local warlords, essentially creating an enabling environment for piracy to thrive in Central Somalia and Puntland.¹ Piracy expanded till 2011 with a total of 246 reported incidents in that year (Figure 1). Within a few months, the Indian Ocean had overtaken other regions as the world's most dangerous for seafarers.

Figure 1: Pirate attacks off Somalia and in the World

Note: The figure plots the number of attacks off Somalia and other regions with piracy problems.
Data source: International Maritime Bureau.

The collapse of piracy after 2011 was as sudden as its rise. IMB attributes 15 piracy attacks to Somali pirates in 2013, 11 in 2014, and none in the first two quarters of 2015 (International Maritime Bureau 2015). As of June 2015, a total of 1,099 attacks on the high seas have been attributed to Somali pirates. Over the 2005-2012 period, pirates of Somalia attacked vessels as far south as the Mozambique Channel and went north all the way to the strait of Hormuz; their catchment area moreover extended west to the south of the

¹The capital region of Mogadishu and southern Somalia were however the subject of attacks from the insurgent group Al-Shabaab, originally the armed wing of the then-defunct ICU. Few pirate attacks overall originated from the pirate stronghold of Kismayo in Southern Somalia.

Red Sea, and on the east to the western shores of the Indian subcontinent. Each dot on the map in Figure 2 represents a reported incident attributed to Somali pirates.

Out of these 1,099 attacks, 216 were successful, resulting in the hijacking of ships and crews (these instances are depicted by a red dot in Figure 2). Figure 3 plots Somali attacks and the success rate, i.e. the fraction of attacks that were successful in a given year. The success rate spiked in the early years of piracy to drop equally sharply in 2009, at the same time that international navies started to get deployed to police the Indian Ocean, and again in 2011 as private armed security contractors were increasingly placed on board of vessels sailing through pirate-infested waters.²

Somali pirates' modus operandi At the origin of a pirate attack is the initiative of an instigator who would assemble a team consisting of an assault crew and a "hold-out" team in charge of logistics once a vessel is brought to shore for ransom negotiations. An assault is conducted out of two to three skiffs with between 10 to 20 persons on board. Pirates would be armed with weapons, be equipped with navigation equipment, and dispose of a ladder to board vessels (World Bank 2013). The targeting of vessels is premeditated rather than the outcome of a random encounter at sea; pirates are all equipped with marine radars, which identify the characteristics (size and type) of vessels with Automatic Identification Systems (AIS) — a mandatory equipment for all large commercial vessels (300 GRT and above) and all passenger ships (International Maritime Organization 2016) and widely adopted otherwise for anti-collision and search-and-rescue purposes.

If a vessel is successfully hijacked, she is brought back to Somali waters and anchored off the coast. A "hold-out team", which comprises a negotiator and a posse dedicated to securing access to the beach to deliver protection, food, water, and energy to the hijacked ship and her crew, then takes over. The negotiation team's connections with local power brokers and the stability of the local political landscape therefore play a critical role on the pirates' ability to sustain protracted negotiations with the shipping companies' representatives (World Bank 2013).

Once an agreement is reached, the ransom is paid and the vessel released. We estimate

²Since 2015, two attacks have been reported but none resulted in a ransom payment.

Figure 2: Attack and Hijack Locations

Note: Each dot indicates the location of an attack attributed to Somali Pirates between 2005 and 2012. The red dots are the successful attacks that resulted in a hijack.

Data Source: International Maritime Bureau.

Figure 3: Number of attacks and success rate, Somalia

Note: The figure plots the number of attacks and the average success rate of attacks, defined as the number of hijacks over the number of attacks in a given year.

Data source: International Maritime Bureau.

that Somali pirates extracted around US\$338 million in ransoms over the 2005-2015 period (see Section 4 for a description of the data). The data also shows that the amount of ransoms extracted and the duration of captivity increased over the years (see Figure 4). In 2012, the data is however censored as some vessels were still being held hostage by 2015.

Figures 4 and 5 highlight two stylized facts about ransoms and durations of captivity. First, payments increased over time and so did hold-out periods for hijacked vessels.³ Second, overall, ransom amounts are positively correlated with durations of captivity (Figure 5 panel a) and the size of the vessel hijacked (Figure 5 panel b).

³The longest reported case is the crew of the MV Albedo, hijacked in 2010 and released only in June 2014 after 1,288 days of captivity; the vessel had sunk in 2013 and among the crew, 5 members had died or were reported missing.

Figure 4: Ransom paid and duration of captivity

Note: The figure plots the annual average ransom paid (in million dollars) and the duration of captivity (in days) of hijacked vessels.

Data source: UNODC-WB.

Figure 5: Ransom, Duration of Captivity, and Vessel Size

Note: The figure plots the log of ransom against the log of duration of captivity and the size of the vessel of each negotiation. The date associated with each point is the year of attack.

Data source: UNODC-WB.

International response In 2008, the U.N. Security Council passed a series of resolutions paving the way for military interventions both within Somalia's territorial waters (U.N. resolution 1816) and onshore (U.N. resolution 1851). Policing off the coast of Somalia was primarily conducted by three international coalitions: U.S.-led Combined Task Force 151 (CTF 151), NATO's *Operation Ocean Shield*, and the European Union Naval Force's *Operation Atalanta*. The military assets started to be deployed in late 2008-early 2009. Some countries such as China or India also sent independent missions to the area. As a result, between 21 and 30 vessels were patrolling the waters off the coast of Somalia at any point in time in 2012; meanwhile 1,190 pirates were being held in custody either serving time or awaiting trial (Oceans Beyond Piracy 2013).

The shipping industry too adopted its own counter-piracy measures. MSCHOA (2011) issued "best management practices for protection against Somalia-based piracy". These include instructions to protect ships against boarding, procedures to follow in case of an attack, and reporting protocols. On the question of private protection, however, no recommendation was provided: "[t]he use, or not, of armed Private Maritime Security Contractors on board merchant vessels is a matter for individual ship operators to decide following their own voyage risk assessment and approval of respective Flag States. This advice does not constitute a recommendation or an endorsement of the general use of armed Private Maritime Security Contractors" (MSCHOA 2011).

Before the International Maritime Organization (IMO) released its *Interim Guidance to Private Maritime Security Companies Providing Privately Contracted Armed Security Personnel On Board Ships in the High Risk Area* in May 2012 (International Maritime Organization 2012), few standards or guidance were available to the shipping industry, therefore hampering the systematic deployment of private maritime security contractors. Article 92 of the U.N. Convention on the Law of the Sea gives each flag State "exclusive jurisdiction on the high seas." However, countries where armed security companies are registered, operate, or transit might also exercise jurisdiction. The surge in piracy off the coast of Somalia has however prompted the international community and individual States to reconsider their legal stances on private armed guards on board of vessels. The U.K. allowed armed security on board towards the end of 2011; France did so only in late 2013.

3 A Model of Crime and Security on the High Seas

We now turn to laying out a model, which incorporates the main features of the piracy sector in Somalia as described in the previous Section while delivering predictions consistent with the data.

Each team of pirates i consists of one assault team and one hold-out team placed under the command of a kingpin. We consider a continuum of teams of measure N that live indefinitely and are characterized by discount rate δ_i and productivity a_i^t . In the first period, $t = 0$, a pirate team i is thus characterized by a pair $\{\delta_i, a_i^0\} \in [\underline{\delta}, \bar{\delta}] \times [b, +\infty)$ and these initial conditions are independently distributed with distributions $F(\cdot)$ and $G^0(\cdot)$, uniform over some interval $[\underline{\delta}, \bar{\delta}]$ and Pareto with scale b and shape θ , respectively.

3.1 Entry, attack, and success rates

At the beginning of each period t , pirates decide whether to enter. Those who do pay a fixed cost f^t .

Vessel protection and pirate success rates To capture the distinctive role of navy patrols and private maritime security contractors, we follow Lazear (2004) and Eeckhout et al. (2010) and consider a policing model in which law enforcement can be applied uniformly or be targeted at one group of vessels only. Uniform law enforcement is ensured by navies from the military coalitions operating in the Indian Ocean. Policing can alternatively be targeted by dispatching security personnel on board selected vessels.

Upon entry, therefore, the probability of successfully hijacking a vessel of size S is a function of vessel size and pirate team ability (hereafter denoted as pirate ability), but also the number of naval assets ν^t deployed to police the Indian Ocean and whether the vessel has private maritime security contractors. As discussed in the previous section, presence of security personnel is only observed by pirates once an attack is actually launched; we denote $\kappa^t(S)$ as the fraction of vessels of size S that have private protection in period t . The level of protection and its composition $\{\nu^t, \kappa^t(\cdot)\}$ are publicly known.

To capture how pirate ability, vessel size, and the presence of international navies

affect hijacking probabilities, we note that larger vessels are intrinsically more difficult to hijack due to, among other things, a higher freeboard, but assume that the ability of assault teams mitigates these obstacles. Furthermore, we consider international navies as being public goods.⁴ We specify the probability of a successful hijacking as a function of pirate ability a , vessel size S , and law enforcement $\{\nu^t, \kappa^t\}$

$$\Pi(S, a, \nu^t, \kappa^t) = [1 - \kappa^t(S)] \cdot e^{-\eta\nu^t} \cdot \pi(S, a). \quad (1)$$

Equation (1) reflects a success probability that is log-separable in each of its components: the fraction $\kappa^t(S)$ of vessels of size S that are protected by private maritime security contractors, the intensity ν^t of navy deployments, ship-specific factors and a pirate-ability term that interact in function $\pi(\cdot)$. For the purpose of the estimation, we specify these functions as follows:

$$\kappa^t(S) = \frac{k^t}{1 + e^{-z(S-S_0)}}, \quad (2)$$

$$\pi(S, a) = \pi_0 \cdot e^{-\zeta \frac{S}{a^\phi}}. \quad (3)$$

The logistic functional form for $\kappa^t(S)$ aims at capturing the feature that, consistent with multiple reports, only the largest vessels are equipped with private maritime security contractors. The “intrinsic” success probability $\pi(S, a)$ allows individual pirate ability to mitigate the effect of vessel size.

Pirate learning by doing and learning from others Finally, as discussed in Section 2, pirates’ success rates are increasing over time, which we attribute to learning. Specifically, we allow for two forms of learning: learning by doing and learning from others.

On the one hand, pirates learn from their own experience (Foster and Rosenzweig 1995) so that their ability is assumed to follow rule of motion

$$a_i^{t+1} = e_i^t \cdot \Lambda(S_i^t, a_i^t) + (1 - e_i^t) \cdot a_i^t, \quad (4)$$

⁴See Fu and Wolpin (2014) for a recent example where crime crowds out police, leading to potential multiplicity of (crime, police) equilibria.

where S_i^t is the size of the vessel attacked in period t and a_i^t is pirate i 's ability at the beginning of period t , and $e^t \in \{0, 1\}$ is the dummy variable that takes value 1 if and only if pirate i entered in period t . We further specify function $\Lambda(\cdot)$ as

$$\Lambda(S, a) = (\lambda \cdot S^\rho + a^\gamma)^{\frac{1}{\gamma}} \quad (5)$$

and impose the following parameter constraint:

A1: Concave learning curve $\Lambda(\cdot)$ is concave in each of its arguments, i.e. $\gamma \in (0, 1)$ and $\rho \in (0, 1)$. ■

On the other hand, pirates learn from others, whereby the past number of attacks makes entry easier. One could think of transfer of know-how from old to new pirates (Conley and Udry 2010) or diffusion of information on the benefits of piracy.

Denoting the measure of entrants in period t by n^t , we assume that the entry cost in t is given by $f^t = f(n^{t-1})$, where

$$f(n) = \psi_0 \cdot (1 + \psi_1 e^{-\psi_2 n}). \quad (6)$$

3.2 Ransom negotiation: holdout and payment

Once hijacked, pirates then attempt to “sell back” cargo and crew to the shipping company. As suggested by the correlations shown in Figure 5, ransom amounts are determined by vessel size and holdout time. Furthermore, longer holdout time is associated with higher ransoms paid. Ransom negotiation is thus modeled as a split-the-pie bargaining game in which pirate discount rates are private information.⁵

⁵Our modeling approach for this specific sub-game is similar to Ambrus et al. (2018) in that both structurally estimate a bargaining game where delay strategically signals one's patience. In the context of ransom negotiations by the Barbary Corsairs of 16th-18th century Europe however, asymmetric information is found to be on the “buyer's” side as longer delays are associated with *lower* ransoms paid.

While vessel size S is observed, vessel value $V(S)$ is known to players but not to the econometrician. We then assume

$$V(S) = V_0 S^\xi. \quad (7)$$

Moreover, shipping companies (buyers) have public-information discount rate δ_0 .

The discount rate of pirates (sellers) on the other hand is assumed to be driven by (i) their actual discount rate δ_i , but also by (ii) their ability a_i^t and (iii) an exogenous shock $\tilde{\varepsilon}^t$, assumed to be i.i.d across time periods and to follow a uniform distribution between 0 and 1.

We thus specify the effective discount rate $\tilde{\delta}_i^t$ as a function of innate discount rate δ_i , ability a_i^t , and shock $\tilde{\varepsilon}^t$:

$$\tilde{\delta}_i^t = \delta_i (1 - \beta_0 e^{-\beta_1 a_i^t}) \tilde{\varepsilon}^t.$$

Earlier, we discussed how critical it is for pirates to be able to maintain hijacked vessels anchored off the coast of Somalia during protracted negotiations. They thus face a probability of losing their catch given the fragility and volatility of the local political landscape. Such probability is driven by their own ability a_i^t and some random shock $\tilde{\varepsilon}^t$. None of the three parameters (discount rate, ability, and shock) are observed by the shipping company, making the pirates' effective discount rate private information.

Assuming that the negotiation process follows the structure in Admati and Perry (1987) and Cramton (1992), the value of the ransom is the complete-information Rubinstein (1982) outcome where shipping companies and pirates have discount rates δ_0 and $\tilde{\delta}_i^t$, respectively. The amount of the ransom paid is then given by $p(\tilde{\delta}_i^t, S)$, where

$$p(\delta, S) = \frac{1 - \delta_0}{1 - \delta_0 \delta} V(S). \quad (8)$$

In a separating equilibrium, a pirate of type $\tilde{\delta}_i^t$ who announces a type δ is offered a ransom $p(\delta, S)$ after delay $D(\delta)$. The holdout time function $D(\cdot)$ that supports such

separating equilibrium satisfies the incentive-compatibility constraint given by :

$$\tilde{\delta}_i^t \in \arg \max_{\delta} [\tilde{\delta}_i^t]^{D(\delta)} p(\delta, S).$$

Taking the first-order conditions yields

$$D'(\tilde{\delta}_i^t) = -\frac{\delta_0}{(1 - \delta_0 \tilde{\delta}_i^t) \cdot \ln \tilde{\delta}_i^t}.$$

Since $\tilde{\varepsilon}^t$ can take value in $[0, 1]$, the support of the distribution of discount rates has 0 as its lower-bound; the holdout function $D(\tilde{\delta}_i^t)$ is then given by the integral

$$D(\tilde{\delta}_i^t) = -\int_0^{\tilde{\delta}_i^t} \frac{\delta_0}{(1 - \delta_0 \delta) \cdot \ln \delta} d\delta, . \quad (9)$$

3.3 Equilibrium vessel choice and entry decision

A pirate team i 's instantaneous gross payoffs from attacking a vessel of size S are equal to

$$\tilde{u}_i(S, a, \nu^t, \kappa^t, \tilde{\delta}_i^t) = \begin{cases} \tilde{\Omega}_i(a, \tilde{\delta}_i^t) \cdot V(S) & \text{with prob. } \Pi(S, a, \nu^t, \kappa^t) \\ 0 & \text{with prob. } 1 - \Pi(S, a, \nu^t, \kappa^t) \end{cases},$$

where $\Pi(\cdot)$ is defined by (1) and $\tilde{\Omega}_i(a, \delta) = \delta^{D(\delta)} \frac{1-\delta_0}{1-\delta_0 \delta}$. When pirates do not enter, we set their utility to $\tilde{u}_i \equiv 0$. We further make the following assumptions on beliefs about the prevalence of private protection κ^t and navy presence ν^t :

A2: Random walk assumption We assume that $\kappa^t(S)$ and $e^{-\eta \nu^t}$ follow a random walk, i.e. $E_t[\kappa^{t+1}(S)] = \kappa^t(S)$, and $E_t[e^{-\nu^{t+1}}] = e^{-\nu^t}$. ■

Under Assumption A2, current law enforcement best predicts future law enforcement, which implies that for every $\tau \leq t$, $E_\tau \Pi(S, a, \nu^t, \kappa^t) = \Pi(S, a, \nu^\tau, \kappa^\tau)$. We integrate over all shocks $\tilde{\varepsilon}^t$ and define $\Omega_i(a) \equiv \int \tilde{\Omega}_i(a, \tilde{\delta}_i^t) dH(\tilde{\varepsilon}^t)$, so we can write for every $t \geq \tau$, a pirate

team i 's indirect utility

$$E_\tau [\tilde{u}_i(S, a, \nu^t, \kappa^t)] = \Pi(S, a, \nu^\tau, \kappa^\tau) \cdot \Omega_i(a) \cdot V(S) \equiv u_i^\tau(S, a).$$

At the beginning of each period τ , pirates observe the level of law enforcement $\{\nu^\tau, \kappa^\tau\}$, and decide whether or not to enter the business, a choice captured by the variable $e^\tau \in \{0, 1\}$, and upon entry, what vessel size S^τ to target. Such nested decision is then governed by optimization of the following program:

$$W_i^\tau(a_i^\tau, n^{\tau-1}) = \sup_{\{e^t, S^t\}} \sum_{t \geq \tau} \delta_i^{t-\tau} \cdot e^t [u_i^\tau(S^t, a_i^t) - f(n^{t-1})],$$

subject to law of motion (4), i.e. $a_i^{t+1} = e^t \cdot \Lambda(S^t, a_i^t) + (1 - e^t) \cdot a_i^t$.

In each period τ , pirates thus face a Markov decision process with state variables $\{a, n\}$ and control variable S . The problem is reminiscent of the durable good problem studied by Rust (1985) as it comprises a two-step decision process: a decision to enter followed by a target choice in case of entry. In any period $t \geq \tau$, the Bellman equation for this Markov decision process can be written as

$$W_i^\tau(a_i^t, n^{t-1}) = \max \left\{ \delta_i W_i^\tau(a_i^t, n^t); \sup_S [u_i^\tau(S, a_i^t) - f(n^{t-1})] + \delta_i W_i^\tau(\Lambda(S, a_i^t), n^t) \right\}.$$

We first observe that atomistic agents do not internalize their entry or vessel choice decisions on current or future entries. Thus, while entry levels are determined in equilibrium, as far as an individual pirate is concerned, the sequence $\{n^t\}^{t \geq \tau}$ is taken as given.

Note that the value function is a non-decreasing function of a and given assumption A1, entry is an absorbing state: if pirates enter in some period $t \geq \tau$, they enter in all subsequent periods *as long as law enforcement remains unchanged*.

Vessel choice In the entry region, the Euler equation can be written

$$\frac{\partial u_i^\tau(S_i^t, a_i^t)}{\partial S} + \delta_i \frac{\partial \Lambda(S_i^t, a_i^t)}{\partial S} \frac{\partial u_i^\tau(S_i^{t+1}, \Lambda(S_i^t, a_i^t))}{\partial a} = 0.$$

Vessel size choice balances static costs versus dynamic gains: larger vessels are on the one hand harder to successfully hijack but when it is the case, generate higher payoffs: for every triplet (S, a, n) ,

$$\frac{\partial u_i^\tau(S, a)}{\partial S} = \Omega_i(a) \left[\frac{\partial \Pi(S, a, \nu^\tau, \kappa^\tau)}{\partial S} \cdot V(S) + \Pi(S, a, \nu^\tau, \kappa^\tau) \cdot V'(S) \right].$$

In addition, as captured by term $\frac{\partial \Lambda(S_i^t, a_i^t)}{\partial S}$, attempting an attack on a larger boat translates into a learning premium as it also leads to increased ability in subsequent periods. The learning premium can be written for a pair (S, a) :

$$\frac{\partial u_i^\tau(S, a)}{\partial a} = \left[\Pi(S, a, \nu^\tau, \kappa^\tau) \cdot \Omega'_i(a) + \frac{\partial \Pi(S, a, \nu^\tau, \kappa^\tau)}{\partial a} \cdot \Omega_i(a) \right] V(S),$$

and comes from both a higher future probability of success and a stronger bargaining position during ransom negotiations.

Entry Since entry is absorbing, once pirate i decides to enter $T_i^\tau(a_i^\tau)$ periods after date τ , the remaining decision is the vessel size $\{S_i^t\}_{t \geq T_i^\tau(a_i^\tau)}$ to attack. In other words, optimal vessel size choices are independent of the actual date of entry.

Moreover, the decision to enter at some period $T + \tau$ trades the cost of entry off against today's static gains, but also future gains from learning.⁶ Entry thus occurs if and only if

$$\sum_{t \geq 0} \delta_i^t u_i^\tau(S_i^{\tau+t}, a_i^{\tau+t}) \geq \frac{1}{1 - \delta_i} f(n^{\tau+T-1}). \quad (10)$$

This leads to two results. First, if $T_i^\tau(a_i^\tau)$ denotes the number of periods past τ after which

⁶If pirate i decides to enter in period $\tau + T$ for some $T \geq 0$, the payoffs he gets are

$$\sum_{t \geq \tau+T} \delta_i^{t-\tau} [u_i^\tau(S_i^{t-T}, a_i^{t-T}) - f(n^{t-1})] = \sum_{t \geq \tau+T} \delta_i^{t-\tau} u_i^\tau(S_i^{t-T}, a_i^{t-T}) - \sum_{t \geq \tau+T} \delta_i^{t-\tau} f(n^{t-1}),$$

while postponing entry by one additional period would yield

$$\delta_i \left[\sum_{t \geq \tau+T} \delta_i^{t-\tau} [u_i^\tau(S_i^{t-T}, a_i^{t-T}) - f(n^t)] \right] = \delta_i \sum_{t \geq \tau+T} \delta_i^{t-\tau} u_i^\tau(S_i^{t-T}, a_i^{t-T}) + \delta_i^T f(n^{\tau+T-1}) - \sum_{t \geq \tau+T} \delta_i^{t-\tau} f(n^{t-1})$$

Comparing these two payoffs yields inequality (10).

pirate i enters, then $T_i^\tau(a_i^\tau)$ is given by

$$T_i^\tau(a_i^\tau) = \min \left\{ T \geq 0, \sum_{t \geq 0} \delta_i^t u_i^\tau(S_i^{\tau+t}, a_i^{\tau+t}) \geq \frac{1}{1 - \delta_i} f(n^{\tau+T-1}) \right\} \quad (11)$$

Second, high ability pirates are the first to enter, followed by their lower-ability peers: entry in period τ is characterized by a cutoff $\underline{a}(\delta_i, n^{\tau-1}, \nu^\tau, \kappa^\tau)$ such that pirate i enters in period τ if and only if

$$a_i^\tau \geq \underline{a}(\delta_i, n^{\tau-1}, \nu^\tau, \kappa^\tau). \quad (12)$$

The measure of the set of entrants (i.e. the number of observed attacks) in period τ is thus given by

$$n^\tau(n^{\tau-1}, \nu^\tau, \kappa^\tau) = \overline{N} \cdot \int_{\underline{\delta}}^1 \left[\int_{\underline{a}(\delta, n^{\tau-1}, \nu^\tau, \kappa^\tau)}^{+\infty} dG^\tau(a) \right] dF(\delta). \quad (13)$$

4 Estimation

4.1 Data Description

The data set that we use in this paper comes from two main sources. The first consists of data on pirate attacks provided by the International Maritime Bureau's Piracy Reporting Center, thereafter referred to as the IMB data. The Piracy Reporting Center collects information on all self-reported pirate-related incidents around the world.⁷ We classify an attack as carried out by Somali pirates if it took place along the coast of Somalia, the Arab Republic of Egypt, Eritrea, the Islamic Republic of Iran, Iraq, Kenya, Oman, Saudi Arabia, Seychelles, Tanzania, the United Arab Emirates, the Republic of Yemen, or anywhere in the Red Sea, Arabian Sea, Gulf of Aden, or Gulf of Oman. We then drop from the data all attacks that are not carried out by Somali pirates. Moreover World Bank (2013) checked local news outlets for information on attacks that took place along the coast of India and in the strait of Hormuz; four incidents were not attributed to Somali pirates and were dropped accordingly from the data. Attacks in our data span from 2000 to 2012,

⁷Given the self-reported nature of the information, under-reporting could be expected in the earlier years when piracy was not yet salient, potentially leading to an underestimate of the effect of law enforcement.

and include the date, time, location (longitude and latitude) of the attack, and basic characteristics of the vessel being attacked, such as type, gross registered tonnage (GRT), etc. The IMB data broadly classify the outcome of the attack into “attempted”, “boarded”, “detained”, “fired upon”, “hijacked”, “missing”, and “suspicious” categories. We define those attacks that ended in hijacks as “successful” attacks.

The second source is the data on ransoms. Information on ransoms comes from the joint data set compiled by the U.N. Office of Drugs and Crime and the World Bank (thereafter referred to as the UNODC-WB data set). The data are compiled from open-source information such as newspaper and reports from national and international law enforcement agencies. Some observations were obtained from interviews with the law firms directly in charge of ransom negotiations. The data set covers 233 vessels hijacked by Somali pirates between April 2005 and December 2012, and contains information on vessel and incident characteristics such as vessel name and date of hijacking, which allows us to match observations with the IMB database. In addition, the data provides information on the most recent (up to December 2012) status of the vessel (whether released, captive, liberated, or sunk), the amount of ransoms paid and the length of the negotiation.

4.2 Estimation

We jointly estimate the model parameters with the method of simulated moments following the ideas outlined in McFadden (1989). There are 23 structural parameters in the model, which we collectively denote as the Θ vector. Our estimation strategy consists of finding the vector Θ that minimizes the weighted distance between a set of simulated moments and their counterparts in the data. Specifically, Θ is defined as:

$$\hat{\Theta} = \operatorname{argmin}_{\Theta} \left\{ \left[m_n - \frac{1}{J} \tilde{m}_n(\Theta; j) \right]' \widehat{W} \left[m_n - \frac{1}{J} \tilde{m}_n(\Theta; j) \right] \right\}, \quad (14)$$

where $n = 1, 2, \dots, 57$ indexes the moments to be matched, and $j \in \{1, 2, \dots, J\}$ indexes the number of simulations, J being the total number of such simulations. m_n is the n -th moment in the data, and $\tilde{m}_n(\Theta; j)$ is the n th moment in the j th simulation of the model conditional on parameter vector Θ . $\widehat{W}$ is the weighting matrix, which is the inverse of the

variance matrix of the data moments.⁸ We compute this variance matrix by bootstrapping with 1000 repetitions.

We estimate the model parameters using data from 2006 to 2012 for the set of moments for which the model has predictions on, and that are observed in the data.⁹ The first set of moments that we match consists of the rates of successful hijacks over time, determined in the model by equation (1). Even though the identification of the structural parameters is jointly determined by the entire set of targeted moments, the success rates in the model over time are a direct function of the parameters governing the probability of success given by equation (1).

The next set of moments used in the estimation corresponds to the bi-annual average number of pirate attacks and the annual average size of ships attacked, normalized to its 2006 mean. In the model, these moments are obtained from the nested problem that pirates solve in order to decide whether to enter and if so, the size of the vessel to target. The number of pirate attacks at every year is determined by the mass of pirates with ability above the cutoff $\underline{a}(\delta_i, n^{t-1}, \nu^t, \kappa^t)$. Therefore, matching the number of attacks over time allows to discipline the parameters ruling the measure of the set of entrants and the entry cost function outlined in equation (6). As the Euler Equation shows, the decision on vessel size is influenced by the parameters that define the learning curve of pirates: even though it is costlier to attack larger ships, by doing so pirates can raise their ability in subsequent years. Hence, the normalized annual averages of size of ships attacked can help pin down the parameters governing the learning law of motion in equation (5).

We also match the annual averages of the size of ships successfully hijacked (normalized to the average size of ships attacked in 2006), which are determined by the prob-

⁸Altonji and Segal (1996) show that using the optimal weighting matrix introduces significant small sample bias. We follow Blundell, Pistaferri and Preston (2008), and use a Diagonally Weighted Minimum Distance approach for our weighting matrix. Under this approach, W is set to the diagonal elements of the optimal weighting matrix, while the terms outside the main diagonal are set to zero.

⁹The choice to start our estimation in 2006 despite having data going back to 2005 is motivated by the discussion of the political context in Somalia during this period discussed in Section 2: pirates faced disruptions from ICU operations on land until 2006. Ignoring these would then affect the identification of the model's structural parameters. From 2006 onwards, as emphasized earlier, no such interferences took place as the ICU was removed from power and the transitional government did not have any power outside of the capital. Piracy was concentrated in the warlord-dominated Central and Northern part of the country, which enjoyed relative stability.

ability of success parameters, and in particular, by the parameters ruling the effect of private protection outlined in equation (2). The identification of these structural parameters hinges on vessels being unable to signal the presence of security contractors *before* an attack takes place. As discussed in Section 2, the choice of targets is typically done using maritime radars and hence prior to a physical encounter. The scope for signaling is thus limited.

Our model also has predictions on the amount of ransoms paid and length of captivity of hijacked ships that are given by equations (8) and (9). While in the model both outcomes are linked to the parameters governing the discount rates of pirates and vessel companies, ransoms are further determined by the parameters defining the value of the vessel outlined in equation (7). We therefore use the average annual ransoms and years in captivity as additional moments to match in the data.

The remaining moments that we match correspond to the second moments of ransoms, years in captivity, and size of ships attacked and hijacked over time. Since pirates in the model vary in productivity and discount rates, our model has predictive power on the dispersion of these four outcomes. The information extracted from these moments allows us to capture the degree of pirate heterogeneity in the data.

To estimate the model, for each period we first solve the problem of the pirates of whether to enter the business or not, and the vessel size to target, conditional on the past entry of pirates to the business and the protection level announced at the beginning of the period. To do this, we first discretize the initial distribution of pirates. As pirate teams are characterized by the pair $\{\delta_i, a_i^t\}$, pirates' distribution of discount rates is discretized into 20 equally spaced points between $\underline{\delta}$ and $\bar{\delta}$. We also discretize the initial distribution of pirates' productivity into 50 equally logarithmically spaced points between b and $\bar{a}$, where b is the theoretical lower bound of the initial distribution of productivity and is fixed at 1.0. Since the Pareto distribution is unbounded from above, we truncate the upper tail by setting $G^0(\bar{a}|b, \theta) = 0.99$, which makes $\bar{a}$ the top 1 percent of the initial distribution of ability. We compute the expected payoffs for every pair in the grid, and find the expected payoff of potential entrants through spline interpolation and numerical integration based on the trapezoid method. The expected payoff of potential entrants

pins down the number of entrants, n^t , in each year.

We then proceed to simulate the model conditional on the number of entrants. In each year, we randomly draw n^t pirate teams from the distribution $G(\cdot)$. For each team, we first compute its success rate based on its vessel choice, and then we simulate the outcome of the attack. Only pirate teams who successfully hijacked a ship enter the negotiation stage and secure ransoms. At the end of the 7-year simulation, we compute the moments from the simulated data. The entire simulation is repeated $J = 100$ times for every input of Θ in the algorithm that minimizes equation (14), and the simulated moments are computed as the average across the 100 simulations.

The estimated 23 parameters for the benchmark model, along with their asymptotic standard errors, are reported in Table 1. At the estimated parameter values, the productivity of pirate teams in the first period is drawn from a Pareto distribution with scale $b = 1$ and shape $\theta = 2.026$. The initial discount rates of pirates are uniformly drawn from the interval $\underline{\delta} = 0.592$ to $\bar{\delta} = 0.969$. The learning process in the model is determined by the parameters $\rho = 0.002$, $\gamma = 0.865$ and $\lambda = 1.031$. Note that assumption A1 in the model is not binding as the estimated coefficient $\gamma < 1$ and $\rho < 1$, ensuring a concave learning curve. The parameters $\beta_0 = 1.617$ and $\beta_1 = 0.917$ determine the influence of pirate productivity a_i^t on the time-varying discount rates.

At the estimated parameters, the size of the average ship attacked by pirates in 2011 was 3.68 times larger than in 2006, and by 2011 and 2012, 75.4 and 84.5 percent of ships this size had security contractors onboard. At these values, the success rate that a pirate team in the 90th percentile of the productivity distribution had of hijacking a ship of this size dropped from 0.41 in 2010 to 0.10 in 2011 and 0.06 in 2012.

The surge in navy patrols also impacted the success rate of attacks, but more so among ships of smaller size. For instance in 2007, with no navy patrols protecting the waters, the probability of hijacking an average-sized ship by 2006 standards (i.e., ships of size 1.0) for a pirate team with average productivity was 0.57, as compared to 0.27 for a ship twice as large (i.e., ships of size 2.0). The increase in navy vessels (to 65 by 2010) dropped the probability of hijacking the smaller ship to 0.42 and the larger ship to 0.25 –an absolute drop of 15.0 and 7.1 percentage points, respectively.

Parameter	Value	S.E.	Notes
θ	2.026	0.151	Baseline shape parameter of the ability distribution
$\underline{\delta}$	0.592	0.094	Lower bound of pirates' discount rate distribution
$\bar{\delta}$	0.969	0.034	Upper bound of pirates' discount rate distribution
η	0.005	0.002	Elasticity of navy patrols in $\Pi(\cdot)$
k^{2011}	0.755	0.12	Location parameter in $\kappa(\cdot)$ in 2011
k^{2012}	0.846	0.064	Location parameter in $\kappa(\cdot)$ in 2012
z	2.465	0.906	Shape parameter in $\kappa(\cdot)$
S_0	0.347	0.047	Threshold parameter in $\kappa(\cdot)$
ζ	2.176	0.344	Shape parameter in $\pi(\cdot)$
ϕ	1.584	0.132	Ability parameter in $\pi(\cdot)$
π_0	1.205	0.233	Location parameter in $\pi(\cdot)$
λ	1.031	0.286	Parameter in $\Lambda(\cdot)$
ρ	0.002	0.001	Shape parameter in $\Lambda(\cdot)$
γ	0.865	0.11	Shape parameter in $\Lambda(\cdot)$
ψ_0	0.236	0.059	Location parameter in $f(N)$
ψ_1	10.002	2.796	Shape parameter in $f(N)$
ψ_2	0.002	0.002	Shape parameter in $f(N)$
β_0	1.617	0.725	Shape parameter in $\delta(\cdot)$
β_1	0.917	1.477	Shape parameter in $\delta(\cdot)$
δ_0	0.733	0.055	Discount factor of shipping companies
V_0	0.354	0.09	Location parameter in $V(\cdot)$
ξ	1.511	0.171	Shape parameter in $V(\cdot)$
$\bar{N}$	600.254	17.846	Location parameter in N

Table 1: Benchmark Parameters

Note: This table reports the 23 benchmark parameters that minimize the distance in Equation 14. Standard errors come from 200 repetition bootstraps. For more details, see Section 4.

In the estimated model, the probability of success is increasing and concave in the productivity of pirate teams. At the estimated parameters, in 2009, pirate teams in the 25th and 75th percentiles of the productivity distribution have a success rate of hijacking a ship of average size by 2006 standards (i.e., of size 1.0) of 0.169 and 0.461 respectively. These same pirate teams have a 0.029 and 0.22 probability of successfully hijacking a twice as large ship of size 2.0.

4.3 Model Fit

The 57 simulated and data moments are listed in Table 2. Figure 6 illustrates graphically the model fit of the first moments, which correspond to the annual averages of: i) success rates, ii) size of ships attacked, iii) size of ships hijacked, iv) ransom amounts, and v) length of captivity of hijacked vessels, as well as the bi-annual average of number of attacks. Figure 7 plots the model fit of the second moments, which are the yearly standard deviations of the: i) size of ships attacked, ii) size of ships successfully hijacked, iii) ransoms paid, and iv) length of captivity.

Overall, simulations of the estimated model track well the dynamics of the piracy business over time. Since our weighting matrix corresponds to the inverse of the variance matrix of the data moments, the fit of the model is more accurate for moments that are more precisely measured in the data than for noisier ones, but most simulated moments lie within the 95 confidence intervals of their data counterparts.

The simulated and actual number of attacks for the years 2005-2012 are presented in Table 2 and Panel (a) of Figure 6. Our model matches closely the steep increase in the number of attacks during the pre-intervention years and the slowdown once navy vessels began patrolling the area. The model also matches well the drop in the number of attacks following the deployment of private maritime security contractors on board vessels.

Table 2 and Panel (b) in Figure 6 show the evolution of success rates over time in the actual and simulated data. From 2006 to 2008, the rate of successful Somali pirate attacks increased from 15 to 31 percent. In 2009, year when navy patrols began protecting the waters in larger numbers, pirates' success rate dropped to 21 percent. The success rate reached 13 percent by 2011, when large vessels began hiring private maritime security contractors. By 2012, success rates increased again to 21 percent as pirates began hijacking smaller ships, who were less likely to hire private maritime security teams. The 95 percent confidence interval of this moment is however quite wide in the data and the annual simulated success rates of the model lie within it. While variation in the simulated success rates is lower, the dynamics over time are similar: in the model, success rates increased from 20 to 24 percent in the pre-intervention years, and then slowed down to 19.8 percent

Moment	Data	Model	Moment	Data	Model
Success rate 2006	0.15	0.20	Avg. N 2006 to 2007	51	70.01
Success rate 2007	0.19	0.23	Avg. N 2007 to 2008	94	110.45
Success rate 2008	0.31	0.24	Avg. N 2008 to 2009	168	152.36
Success rate 2009	0.21	0.20	Avg. N 2009 to 2010	214	203.34
Success rate 2010	0.22	0.20	Avg. N 2010 to 2011	216	256.72
Success rate 2011	0.13	0.10	Avg. N 2011 to 2012	141	185.95
Success rate 2012	0.21	0.28	SD. ship size attacked 2007	1.35	0.97
Avg. ship size attacked 2007	1.3	1.36	SD. ship size attacked 2008	3.5	1.16
Avg. ship size attacked 2008	2.88	1.93	SD. ship size attacked 2009	2.74	1.84
Avg. ship size attacked 2009	2.64	2.75	SD. ship size attacked 2010	2.9	2.20
Avg. ship size attacked 2010	2.74	3.25	SD. ship size attacked 2011	2.84	2.26
Avg. ship size attacked 2011	3.24	3.68	SD. ship size attacked 2012	3.65	2.21
Avg. ship size attacked 2012	3.62	4.42	SD. ship size hijacked 2007	0.21	0.66
Avg. ship size hijacked 2007	0.27	1.25	SD. ship size hijacked 2008	2.38	1.15
Avg. ship size hijacked 2008	1.5	2.04	SD. ship size hijacked 2009	2.25	1.82
Avg. ship size hijacked 2009	1.43	2.92	SD. ship size hijacked 2010	2.22	2.20
Avg. ship size hijacked 2010	1.3	3.64	SD. ship size hijacked 2011	3.01	2.27
Avg. ship size hijacked 2011	2.21	2.87	SD. ship size hijacked 2012	2.37	1.64
Avg. ship size hijacked 2012	1.75	2.23	SD. ransom 2006	0.18	0.24
Avg. ransom 2006	0.13	0.18	SD. ransom 2007	0.45	0.52
Avg. ransom 2007	1.18	0.60	SD. ransom 2008	0.61	1.14
Avg. ransom 2008	1.57	1.35	SD. ransom 2009	1.71	2.14
Avg. ransom 2009	2.63	2.47	SD. ransom 2010	2.06	2.76
Avg. ransom 2010	4.18	3.46	SD. delay 2006	0.15	0.07
Avg. delay 2006	0.18	0.13	SD. delay 2007	0.44	0.15
Avg. delay 2007	0.31	0.28	SD. delay 2008	0.15	0.20
Avg. delay 2008	0.17	0.39	SD. delay 2009	0.22	0.23
Avg. delay 2009	0.27	0.45	SD. delay 2010	0.27	0.24
Avg. delay 2010	0.51	0.47			

Table 2: Moments in Estimation

Note: The table lists the 57 moments matched in the estimation using the Method of Simulated Moments. The “data” column lists the moments in the data, and the “model” column lists the corresponding simulated moments at the optimized $\hat{\Theta}$. Section 4 discusses the selection of moments and estimation method in detail.

Figure 6: Model fit of first moments

Note: The figures plot the average number of attacks, success rates, size of attacked and hijacked ships, ransoms and length of captivity in the data v.s. the model simulated at the benchmark parameters listed in Table 1. The confidence intervals are 1-standard-deviation bands. Data sources: IMB (2012), UNODC-WB, and authors' calculations.

Figure 7: Model fit of second moments

Note: The figures plot the second moments by year of the size of attacked and hijacked ships, ransoms, and length of captivity in the data v.s. the model simulated at the benchmark parameters listed in Table 1. The confidence intervals are 1-standard-deviation bands.

Data sources: IMB (2012), UNODC-WB, and authors' calculations.

by 2009 and 2010. In 2011, the simulated success rate was 10.2 percent. However by 2012, simulated pirates that were still in the business hijacked smaller ships as a response of private maritime security, increasing success rates to 28.4 percent.

Panels (c) and (d) in Figure 6 plot the average size of ships targeted by pirates and effectively hijacked over time, both in the actual and simulated data. The data shows that the average size of the ships attacked by pirates increased over the years: in 2007, ships attacked were 30 percent larger than the previous year and by 2012 they were 3.6 times larger. With the presence of navy patrols, pirates targeted ships slightly smaller but over

time the size of ships attacked continued growing. Our model captures the positive trend of size of ships attacked over time, and for most years the simulated moments are within the 95 confidence interval of their data counterparts.

Pirates were also able to hijack larger ships over time, a pattern that our model matches, but less precisely as the data moment is noisier. From 2006 to 2012, hijacked ships were 1.75 and 2.22 times larger in the actual and simulated data, respectively. While our model over-predicts the average size of hijacked ships, it captures the reduction in size of hijacked ships in 2012 after the surge of private security.

Table 2 and panels (e) and (f) in Figure 6 compare the actual and simulated means of ransom payments (measured in millions of dollars) and length of captivity (measured in years) over time. Our model predictions on ransoms are close to their data counterparts and show a constant increase in the average ransom obtained by pirates throughout the years. In the data, the variance of length of captivity is substantially high and our model overestimates the length of captivity for the years 2008 and 2009.

5 Simulations

5.1 The role of learning in the rise of Somali piracy

Our model allows for two forms of learning: learning by doing (or private learning) and learning from others (or social learning). Following the law of motion of learning-by-doing, as pirates attack larger ships, their productivity in subsequent periods rises. On the other hand, learning-from-others, in the model, occurs through accumulated experience reflected in the number of past attacks: as more pirates enter the business, the entry costs in future periods drop.

We use the estimated model to understand the extent to which these two forms of learning can explain the surge of pirate attacks during the years preceding the deployment of protective measures. To do this, we simulate two counterfactual scenarios in which: (i) the productivity of pirates stays fixed over the years (i.e., $\lambda = 0$); and (ii) the fixed costs of entry do not vary with the number of past entrants (i.e., $\psi_2 = 0$). These two

counterfactual simulations allow us to separately shut down private and social learning, and examine how the piracy business would have evolved as compared to our benchmark model.

The results from the benchmark model and the two counterfactual exercises are presented in Figure 8. As the red dashed line shows, in the absence of learning-by-doing, it is not profitable for pirates to enter the business in the first place as they do not expect to become better over time. Once the dynamic gains that come from private learning disappear, the remaining expected gains are not high enough to cover the costs of entering. Learning-by-doing is, therefore, critical in our model to explain the entry of pirates to the business. However, learning-by-doing alone cannot explain the steep increase in the number of attacks observed in the years before law enforcement intensified. In a world where there is no learning from others, only the more productive pirates decide to enter the business in the first period. Since in the absence of social learning entry costs do not decline over time, pirates of lower initial productivity are prevented from entering the business in posterior years. By shutting down learning-from-others, the number of attacks from 2006 to 2009 declines from 445 to 195; that is, by 43.8 percent.

The two forms of learning considered here are critical for explaining the rise of Somali piracy during the years 2005-2012. The first pirates to enter are the most able who incur negative payoffs today in exchange of positive payoffs in the future. As they enter, their ability improves and they also lower the entry costs for future cohorts, who in turn learn by doing and further lower entry costs.

5.2 The impact of private protection

We now use the estimated model to examine the effectiveness of having private armed contractors deployed on merchant vessels.

Counterfactual of no private protection. To analyze the role of private maritime security contractors, we simulate a counterfactual world without this type of protection. We do this by setting $k^{2011} = k^{2012} = 0$ throughout the entire simulation, while keeping all the other parameters and variables as in the benchmark model. In particular, the intensity of

Figure 8: Role of learning on the number of attacks

Note: The figure plots the total number of attacks in the benchmark model (blue solid line) and two counterfactual scenarios where we shut down private learning (red dash line) and social learning (red solid line).

navy patrols is kept untouched. The results of the counterfactual analysis are presented in Figure 9.

The first order impact of removing private protection is observed in increased success rates, allowing average ransoms to expand in 2011 and 2012 as seen in Figure 9. Without private protection, success rates hovered around 23 percent in 2011 and 2012 as pirates continue to prey on the largest vessels. By contrast, success rates in the benchmark simulations dropped down to 10 percent by 2011, the year when private maritime security contractors were deployed.

With the most valuable vessels left unprotected, the average ransom in 2012 surges to US\$4.11 million in the counterfactual, more than doubling the average ransom of the benchmark simulation (US\$1.72 million). Expected future profits increase in the absence of private maritime security, indirectly attracting new pirates into the business. Upon observing the absence of armed contractors on board vessels in 2011, pirates expect the same in 2012 as well. As a result, many potential pirates join the business: the total number of attacks skyrockets to 417 by 2012 in the counterfactual, instead of dropping to

Figure 9: Counterfactual: No private maritime security

Note: The figures plot the entry costs, average pirate ability, total number of attacks, average ransoms, success rates, average size attacked, average size hijacked, and relative delay in the counterfactual analysis where we shut down private maritime security (dash line) and benchmark model (solid line). All other parameters in the counterfactual simulations are kept the same as in the benchmark model.

90 in the benchmark.

Interestingly enough, the removal of security details on board large vessels increases attacks among *all* vessel sizes. Private protection generates a deterrence externality: the removal of private maritime security contractors attracts more potential pirates into the business as the net present value of entering has now increased. Inexperienced new pirates enter the business to prey on smaller vessels with an expectation of learning-by-doing whereas the most experienced ones target larger ships. This externality is similar to the one identified by Ayres and Levitt (1998), and is present in our simulations: the average size of ships attacked in the counterfactual case is lower by around 30 percent relative to the baseline in 2012. The drop in size is due to a change in the composition of pirates: a lower incidence of armed vessels allows more hence less experienced pirates to stay in the business. Consequently, the average size of attacked vessels drops. Within the subset of pirates that survived in the baseline simulation, the average size targeted *increased* by 41 percent once private protection is removed.

The crime elasticity of private protection. To measure the elasticity of crime with respect to the prevalence of private maritime security contractors, we simulate the model several times varying the values of k^{2011} from 0.5 to 1. This exercise allows us to capture changes in the intensity of private maritime security in 2011, while keeping all other parameters constant. The results are reported in Figure 10.

The elasticity of the total number of attacks with respect to private maritime protection is high at around 35.3. As k^{2011} increases from 0.5 to 0.77, the number of attacks in 2012 drops from 282 to 0. The number of attacks stays at zero for k^{2011} greater than 0.77. To obtain the elasticity, we first convert k^{2011} into the percentage of vessels protected. In our simulations, 32.3 and 49.2 percent of all vessels are protected when $k^{2011} = 0.5$ and 0.75, respectively. Around the baseline estimate of $k^{2011} = 0.75$, a one-percentage-point increase in the share of vessels protected reduces the number of attacks by 35.3 percent, and *increases* the number of hijacks by 21.2 percent. The increase in the number of hijacks is due to the strong displacement effects discussed in detail below.

The estimated elasticity means that approximately one pirate attack in 2012 is elim-

Figure 10: Marginal effects of armed private maritime security

Note: The graph plots the marginal effects of $\kappa^{2011}(S)$ on the number of attacks and hijacks. The value of k^{2011} varies between 0.5 and 1, while all other parameters of the model are set to their benchmark value. We have omitted the values of k^{2011} greater than 0.8 because the number of attacks and hijacks are all zeros. The dashed vertical line indicates the baseline value of $k^{2011} = 0.75$.

inated for every 4.6 private maritime security teams hired in 2011. To arrive at this interpretation, we carry out a simple calculation using the Suez canal traffic as a rough proxy for the traffic volume through the piracy-affected waters. In 2012, 14,765 large ships passed through the Suez canal.¹⁰ A one percentage point increase in private maritime security hires (148 hires) leads to 32 fewer attacks (35.3 percent of the 91 attacks in 2012), which in turn implies that $148/32 \approx 4.6$ private teams can reduce the number of attacks by 1 through the mechanisms outlined in our model. On a side note, the cost of 4.6 security teams, as estimated by (Oceans Beyond Piracy 2015), is between \$92,000 and \$220,800.

Even though private maritime security reduces the number of pirate attacks, it also results in a significant displacement externality towards smaller ships. Figure 10 highlights this effect by showing the number of attacks and hijacks to ships below (red dashed) and above (black dash-dotted line) the median size. Around the baseline estimate, a one-percentage-point increase in protection lowers the number of attacks (hijacks) on large vessels by 42.9 (43.8) percent, and at the same time, increases the number of attacks (hi-

¹⁰We classify the following vessel types in the published traffic data as “large”: tankers, LNG ships, bulk carriers, combined carriers, container ships, and car carriers.

jacks) on small vessels by 22.1 (50.6) percent. The positive elasticity indicates that around the baseline level, the displacement externality dominates the deterrence externality, making smaller vessels more vulnerable to the selective provision of security. The displacement effects are particularly strong in the number of hijackings because smaller vessels are also easier to hijack. Nevertheless, as private maritime security intensifies, the deterrence externality eventually dominates the displacement effect. At values of k^{2011} above 0.77 in our baseline estimation, all pirates are pushed out of business.

Since social benefits and distributional implications of private protection are not internalized by ship-owners, there is a potential for regulating the private security market and defining the role of navy patrols, to which we now turn.

5.3 The impact of navy patrols

We study the effects of navy patrols in a similar fashion. In the data and our benchmark simulation, policing troubled waters started in 2009 with 48 patrols. By 2012, the number of patrols had increased steadily to 61. To understand the impact of navy patrols, we first simulate a counterfactual scenario where the number of navy patrols is set to zero for all the years, while keeping all other parameters at their benchmark values. We then estimate the crime elasticity by simulating the model varying the number of navy patrols. The results are summarized in Figures 11 and 12.

Counterfactual of no navy patrols. Withdrawing navy patrols directly improves the odds of successful attacks. In our counterfactual simulation, success rates increase to 24 and 25 percent by 2009 and 2010, respectively. The 2012 success rate climbs to 50 percent, significantly higher than in the data (21 percent) and benchmark model (28 percent).

Different from private maritime security, where average ransoms surge up in its absence, the average ransom is lower than the benchmark case in the case of no-navy. The divergence is caused by the fact that the removal of the navy in itself does not distort the selection of vessels, whereas the removal of private maritime security teams attracts pirates into larger and more valuable ships. For the same reason, the average duration of captivity exhibits a similar effect.

Figure 11: Counterfactual: No navy patrols

Note: The figures plot the entry costs, average ability, total number of attacks, average ransoms, success rates, average size attacked, the average size hijacked, and relative delay in the counterfactual analysis where we shut down navy patrols (dash line) and benchmark model (solid line). All other parameters are kept the same as in the benchmark model.

As with private security, removing the navies increases the expected profitability of piracy, and thus attracts new pirates into the business. As a result, the number of attacks increases to 183 and 306 by 2009 and 2010, respectively. The surge in number of attacks stopped only after private protection was introduced in 2011 in the counterfactual simulations: by the end of 2012, the number of attacks slightly dropped to 307 (still significantly higher than the 90 attacks under the benchmark simulation).

New pirates target on average smaller vessels, as shown by lower average size of ships attacked and hijacked, and lower average ransoms. As with private maritime security, the shift towards smaller ships reflects a preference for smaller vessels among new entrants, which is due to inexperience. In this particular counterfactual, the bias towards smaller vessels is further exacerbated in 2011 as larger vessels start hiring private maritime security contractors in 2011.

The crime elasticity of navy deployment. To put a rough estimate on the elasticity of crime with respect to navy patrols, we run another set of counterfactual analysis in which we gradually double the number of patrols in 2009 from 48 in the benchmark to 96. We report the number of attacks and hijacks in the first two panels of Figure 12.

The crime elasticity with respect to navy deployment is an order of magnitude lower than that with respect to private maritime security. A 1 percent increase in the number of patrols around the benchmark value of 48 reduces the number of attacks (hijacks) by around 0.26 (0.19) percent. Even with the navy patrols doubling to 96, the number of attacks only decreases by 24.6 percent from 232 to 175 in the next year, and the number of hijacks decreases by 17.4 percent from 46 to 38. Compared with the crime elasticity of private protection at 35.3 for the number of attacks, naval deployments are much less effective in either measure. The relative effectiveness of private security is mainly because it is targeted towards the largest vessels, where pirates expect to generate the bulk of their income. Once the most valuable vessels are effectively protected, pirates are less motivated to join the business in the first place due to lowered expected return. On the other hand, navy patrols provide moderate and uniform protection over all the vessels. The uniformity of the protection also implies that large vessels are not as effectively pro-

Figure 12: Marginal effect of navy patrols

Note: the figures plot the marginal effects of navy patrols in year t on the number of attacks and the number of hijacks in year $t + 1$. The value of navy patrols varies between 48, the original value from the data, to 96 in the first two panels, and varies between 0 and 100 in the others.

tected as in the case of private security, which in turn incentivizes pirates to stay in the business.

Naval interventions in the earlier years are more effective in curbing the piracy problem. The next four panels of Figure 12 repeat the previous exercise with navy patrols varying between 0 and 100 in 2007 and 2011. In the earlier stage of piracy in 2007, 20 naval patrols per year are able to completely eradicate the piracy problem; meanwhile the same 20 naval patrols in 2011 can only slightly suppress the number of attacks from around 220 to 180, and the number of hijacks from around 85 to 70. The deteriorating effects of this law enforcement intervention are mainly due to learning-by-doing: the average ability of pirates in 2011 is 2.4 times higher than that in 2007, making them much better suited to survive such policy interventions.

5.4 Optimal allocation of anti-piracy resources

The above counterfactual simulations show that navy patrols and private maritime security work differently in curbing pirate attacks. We now turn to study the cost-effective combination of navy patrols and security teams to tackle the piracy problem. We show that while private maritime security teams are more cost-effective in curbing piracy, they cannot fully eradicate the problem due to displacement effects. Concurrent navy patrols are a necessary condition to eliminate the pirates of Somalia.

The starting point of our analysis is the construction of “iso-crime” curves (Figure 13) of the two policies: each curve plots the combinations of patrols and private protection that will result in the same number of attacks in the following year. The iso-crime curves highlight the advantages and the limitations of each policy. At lower intensities of both navy patrols and private security, the iso-crime curves are close to straight lines as in the first three panels of Figure 13, indicating that public and private provision of security are close substitutes. However, as suggested in the last panel of the same figure, the iso-crime curves become convex at higher levels of policy measures. In particular, the iso-crime curve is almost flat for higher levels of private security, mainly due to the strong displacement effects toward smaller vessels: as a higher fraction of large ships is

protected, pirates are more likely to switch to less-protected smaller ships. The convexity of the iso-crime curves hints that it is hard to eliminate the piracy problem using a single policy measure, and that a cost-effective policy would consist of a combination of both public and private protection instruments.

To estimate the cost-effective combination of policy interventions to reduce the number of attacks, we first estimate the costs of private and public security provision. The cost estimates of private maritime security contractors are available in the literature. While studies such as Besley, Fetzer and Mueller (2015) or World Bank (2013) give aggregate costs associated with Somali piracy, they do not provide a breakdown of these costs. The methods used are based on increased private costs borne by shipping companies and are also likely to ignore the contributions of taxpayers of each country involved in the naval coalitions. The One Earth Future Foundation, a not-for-profit organization, first provided the estimates of private maritime security in their Oceans Beyond Piracy (OBP) 2012 report, and they set the cost of private maritime security contractors at US\$50,000 per trip through the troubled waters (Oceans Beyond Piracy 2013). In the later years, they provided more detailed estimates: at around US\$33,250 per trip in 2013 (Oceans Beyond Piracy 2014), and US\$22,975 per trip in 2014 (Oceans Beyond Piracy 2015). For our exercise, we use the year-specific estimates of costs after 2011, and the rough estimate of US\$50,000 per trip for earlier years. To map the per trip cost to an increase in k^t , we also need information on the traffic volume through the troubled waters. We again turn to the traffic statistics from the Suez canal for a traffic estimate in each year, and compute the costs of a one percentage point increase of protected vessels per year. As the Suez canal traffic underestimates overall traffic through the troubled waters off Somalia, our estimates of private team costs tend to be underestimations as well. For this reason, we carry out a set of robustness checks with doubled costs of armed protection. The results presented in Figure 16 are essentially the same as in our benchmark case.

The estimated costs of navy patrols also come from various OBP reports. In 2010, OBP provided a first estimate of around US\$30.2 million per patrol – including the costs of personnel, fuel, and other administrative supports (Oceans Beyond Piracy 2013). This estimate is almost surely an upper bound, as it assumes that the warship patrols year-around.

Figure 13: Navy patrols vs. private maritime security: iso-crime curves for all vessels

Note: The figures plot the combination of navy patrols and intensity of private maritime security in year t that achieves a certain number of attacks on all vessels in year $t + 1$. The black line in each graph indicates the optimal combination of navy patrols and private maritime security contractors that minimizes the total costs for each level of attacks. All other model parameters are set to their benchmark value.

Figure 14: iso-crime curves for small and large vessels

Note: the figures plot the combination of navy patrols and the strength of private maritime security in year t that achieves a certain number of attacks on large and small vessels in year $t + 1$. The black line in each graph indicates the optimal combination of navy patrols and private maritime security contractors that minimizes the total costs for each level of attacks. All the other parameters are set to their benchmark value.

In the following years, the OBP reports provided detailed breakdown of the costs of patrols. The 2012 report estimated the total costs of operation to be around US\$960 million for 63 patrols, and for the next year, US\$912 million for 61 patrols. We take the difference of these two years, and estimate the cost per patrol to be around US\$23.7 million. This cost estimate is closer to the lower bound of the true costs, as the difference-estimator takes away any fixed costs of maintaining navy bases in the region. In the end, we use the average of the two estimates, US\$27.0 million per patrol for all the years. With the costs of both measures estimated, we then proceed to compute the cost-minimizing combination of policy measures for every level of attack-reduction in each year. We plot the optimal policy rules as the black solid lines in Figure 13.

Private armed protection appears relatively cost-effective in curbing the number of attacks at the early stage. In early years, private protection alone is able to eradicate the piracy problem. In later years, it is also more cost-effective to first implement private protection of largest vessels to reduce the number of attacks from around 300 attacks per year to fewer than 100 attacks. However, private protection of largest vessels alone cannot entirely solve the piracy problem due to crime displacement. Displacement effects

are particularly pronounced after 2011, when further increasing private protection would only encourage pirates to shift toward smaller vessels while leaving the total number of attacks unaffected. At these levels of private security, introducing naval patrol is a necessary condition to effectively deter pirates.

To further highlight the displacement effect, we split ships by size in 2012 and construct the iso-crime curves for small (below median) and large (above median) ships separately (Figure 14). The vertical curves for large vessels imply that private protection is the only effective policy to protect large vessels. Unsurprisingly, the optimal policy for protecting large vessels is to rely exclusively on private maritime security contractors. However, for smaller vessels, private protection leads to more attacks due to displacement, therefore the optimal policy to protect smaller ships is to *reduce* armed guards and at the same time, deploy more navy patrols. The role of the navy then becomes instrumental in keeping the number of attacks on small ships low.

These findings highlight the distributional implications of various policy instruments, which might matter when determining what the optimal level and composition of policing should be.

6 Conclusion

We provide quantitative evidence on the relative contribution that navy patrols and private maritime security contractors had in explaining the collapse of Somali piracy in 2012. The critical role of private security was highlighted, with emphasis made on the externalities – both positive and negative – generated by private measures. The analysis brings new mechanisms to consider when exploring the tension between private and public provision of security.

References

Admati, Anat R. and Motty Perry, “Strategic Delay in Bargaining,” *The Review of Economic Studies*, July 1987, 54 (3), 345–364.

- Altonji, Joseph G and Lewis M Segal**, "Small-sample bias in GMM estimation of covariance structures," *Journal of Business & Economic Statistics*, 1996, 14 (3), 353–366.
- Ambrus, Attila, Eric Chaney, and Igor Salitskiy**, "Pirates of the Mediterranean: An empirical investigation of bargaining with asymmetric information," *Quantitative Economics*, 2018, 9 (1), 217–246.
- Ayres, Ian and Steven D. Levitt**, "Measuring Positive Externalities from Unobservable Victim Precaution: An Empirical Analysis of Lojack," *The Quarterly Journal of Economics*, 1998, 113 (1), pp. 43–77.
- Banerjee, Abhijit, Esther Duflo, Daniel Keniston, and Nina Singh**, "Crime, Punishment, and Monitoring: Deterring Drunken Driving in India," Technical Report, Manuscript 2012.
- Becker, Gary S.**, "Crime and Punishment: An Economic Approach," *Journal of Political Economy*, 1968, 76 (2), 169–217.
- Besley, Timothy, Thiemo Fetzer, and Hannes Mueller**, "The Welfare Cost of Lawlessness: Evidence from Somali Piracy," *Journal of the European Economic Association*, 2015, 13 (2), 203–239.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston**, "Consumption inequality and partial insurance," *The American Economic Review*, 2008, pp. 1887–1921.
- Clotfelter, Charles T.**, "Public Services, Private Substitutes, and the Demand for Protection against Crime," *The American Economic Review*, 1977, 67 (5), 867–877.
- , "Private security and the public safety," *Journal of Urban Economics*, 1978, 5 (3), 388–402.
- Conley, Timothy G. and Christopher R. Udry**, "Learning about a New Technology: Pineapple in Ghana," *American Economic Review*, March 2010, 100 (1), 35–69.
- Corman, Hope and H. Naci Mocan**, "A Time-Series Analysis of Crime, Deterrence, and Drug Abuse in New York City," *American Economic Review*, 2000, 90 (3), 584–604.

- Cornish, Derek and Ronald Clark**, "Understanding Crime Displacement: An Application of Rational Choice Theory," *Criminology*, November 1987, 25 (4), 933–947.
- Cramton, Peter C.**, "Strategic Delay in Bargaining with Two-Sided Uncertainty," *The Review of Economic Studies*, January 1992, 59 (1), 205–225.
- Di Tella, Rafael and Ernesto Schargrodsy**, "Do Police Reduce Crime? Estimates Using the Allocation of Police Forces after a Terrorist Attack," *The American Economic Review*, 2004, 94 (1), pp. 115–133.
- , **Sebastian Galiani, and Ernesto Schargrodsy**, "Crime Distribution and Victim Behavior during a Crime Wave," in Rafael Di Tella, Sebastian Edwards, and Ernesto Schargrodsy, eds., *The Economics of Crime: Lessons for and from Latin America*, University of Chicago Press, 2010, pp. 175–204.
- Draca, Mirko and Stephen Machin**, "Crime and Economic Incentives," *Annual Review of Economics*, 2015, 7, 389–408.
- , —, **and Robert Witt**, "Panic on the Streets of London: Police, Crime, and the July 2005 Terror Attacks," *American Economic Review*, August 2011, 101 (5), 2157–81.
- Eeckhout, Jan, Nicola Persico, and Petra E. Todd**, "A Theory of Optimal Random Crackdowns," *American Economic Review*, June 2010, 100 (3), 1104–35.
- Ehrlich, Isaac**, "Participation in Illegitimate Activities: A Theoretical and Empirical Investigation," *Journal of Political Economy*, 1973, 81 (3), 521–65.
- Foster, Andrew D. and Mark R. Rosenzweig**, "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture," *Journal of Political Economy*, 1995, 103 (6), 1176–1209.
- Fu, Chao and Kenneth I. Wolpin**, "Structural Estimation of a Becker-Ehrlich Equilibrium Model of Crime: Allocating Police Across Cities to Reduce Crime," RISE Working Paper 14-020, Rice University August 2014.

- Galiani, Sebastian, Ivan Lopez Cruz, and Gustavo Torrens**, “Stirring Up a Hornets’ Nest: Geographic Distribution of Crime,” *Journal of Economic Behavior & Organization*, 2018, 152, 17–35.
- Gonzalez-Navarro, Marco**, “Deterrence and Geographical Externalities in Auto Theft,” *American Economic Journal: Applied Economics*, 2013, 5 (4), 92–110.
- International Maritime Bureau**, “Piracy and Armed Robbery against Ships: Report for the period 1 January - 31 December 2005,” Technical Report, International Chamber of Commerce - International Maritime Bureau 2006.
- , “Piracy and Armed Robbery against Ships: Report for the period 1 January - 30 June 2015,” Technical Report, International Chamber of Commerce - International Maritime Bureau 2015.
- International Maritime Organization**, “Interim Guidance to Private Maritime Security Companies Providing Privately Contracted Armed Security Personnel on Board Ships in the High Risk Area,” 2012.
- , “Regulation for carriage of AIS,” 2016.
- Kremer, Michael and Jack Willis**, “Guns, Latrines and Land Reform: Dynamic Pigouvian Taxation,” *American Economic Review*, 2016, 106 (5 (May 2016)), 83–88.
- Lazear, Edward P.**, “Speeding, Tax Fraud, and Teaching to the Test,” NBER Working Papers 10932, National Bureau of Economic Research, Inc November 2004.
- Levitt, Steven D.**, “Using Electoral Cycles in Police Hiring to Estimate the Effect of Police on Crime,” *The American Economic Review*, 1997, 87 (3), pp. 270–290.
- McFadden, Daniel**, “A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration,” *Econometrica*, September 1989, 57 (5), 995–1026.
- MSCHOA**, *Best Management Practices for Protection against Somalia Based Piracy*, Edinburgh, U.K.: Witherby Publishing Group Ltd, 2011.

- Oceans Beyond Piracy**, "The Economic Cost of Somali Piracy 2012," Technical Report, One Earth Future Foundation 2013.
- , "The Economic Cost of Somali Piracy 2013," Technical Report, One Earth Future Foundation 2014.
- , "The Economic Cost of Somali Piracy 2014," Technical Report, One Earth Future Foundation 2015.
- Rubinstein, Ariel**, "Perfect Equilibrium in a Bargaining Model," *Econometrica*, January 1982, 50 (1), 97–109.
- Rust, John**, "Stationary Equilibrium in a Market for Durable Assets," *Econometrica*, July 1985, 53 (4), 783–805.
- van Ours, Jan C. and Ben Vollaard**, "The Engine Immobiliser: A Non-starter for Car Thieves," *The Economic Journal*, 2015, pp. 1264–1291.
- Vollaard, Ben and Jan C. van Ours**, "Does Regulation of Built-in Security Reduce Crime? Evidence from a Natural Experiment*," *The Economic Journal*, 2011, 121 (552), 485–504.
- World Bank**, *The Pirates of Somalia: Ending the Threat, Rebuilding a Nation*, Washington DC, USA: The World Bank, April 2013.

A Additional Figures and Tables

Figure 15: $\kappa^t(\cdot)$ Function

Note: The figure plots the $\kappa^t(S)$ function with our benchmark estimation. $\kappa^t(S)$ is the probability of a ship with private security as a function of the vessel size, S , in year t .

Figure 16: Robustness check: iso-crime curves for all vessels with doubled costs for private security

Note: the figures plot the combination of navy patrols and the strength of private security in year t that achieves a certain number of attacks on all vessels in year $t + 1$. The black line in each graph indicates the optimal combination of navy patrols and private teams that minimizes the total costs for each level of attacks. All the other parameters are set to their benchmark value.