

The Dynamics of Evasion: The Price Cap on Russian Oil Exports and the Amassing of the Shadow Fleet*

Diego S. Cardoso,^a Stephen W. Salant,^b & Julien Daubanes^c

January 6, 2025

Abstract

To reduce funds for Russia's Ukraine invasion, Western governments imposed a price ceiling on Russian seaborne oil exports using Western services. To sell above that ceiling, Russia developed a "shadow fleet" which uses no such services. We use a calibrated model driven by this fleet's expansion to assess various sanctions. While all sanctions reduce the present value of Russia's profits, we find that the tighter the ceiling and the tighter the enforcement, the less harm sanctions impose, contradicting conventional wisdom based on Hotelling lemma. However, policies to reduce the shadow fleet's size may increase the sanction's effectiveness.

JEL codes: D04, L51, Q41

Keywords: Economic warfare; Sanctions evasion; Hotelling's lemma

*We have benefited from useful comments by participants at various conferences and seminars: Copenhagen International Workshop in Energy and Resource Economics and Policy, Midwest Economic Association, University of Illinois Urbana-Champaign, University of Copenhagen. Particular thanks to Tatyana Deryugina, Gérard Gaudet, Ryan Kellogg, Lutz Kilian, Lukasz Rachel, Daniel Spiro, and Catherine Wolfram. The research leading to these results has been supported by the Swiss National Science Foundation (Grant 192578).

^aDepartment of Agricultural and Consumer Economics, and University of Illinois Urbana-Champaign.

^bDepartment of Economics, University of Michigan, and Resources for the Future (RFF).

Corresponding author: ssalant@umich.edu.

^cDepartment of Technology, Management and Economics, Technical University of Denmark, Center for Energy and Environmental Policy Research, Massachusetts Institute of Technology, and CESifo.

1 Introduction

In response to Russia’s invasion of Ukraine in February 2022, the European Union, the United States, and other G-7 countries ceased their imports of Russian oil, leading Russia to export more to India, Turkey, and China instead. In addition, the G-7 and other allies imposed sanctions on oil exports from Russia in order to limit its ability to finance the war. Russia’s revenue from the sale of crude oil and related petroleum products is instrumental in supporting its government spending. Before the invasion, oil-related revenues amounted to more than 40% of Russia’s federal budget (International Monetary Fund, 2021, p.33). More than 17% of that budget goes to pay for the Ukraine war.¹ Given this situation, RAND concluded: “Blocking those [revenues] could be the most powerful tool in the West’s economic toolkit to hamper Russia’s war effort...” (Shatz & Reach, 2023, p.vi).

Since more than 80% of Russia’s seaborne oil exports relied on the provision of insurance and other Western services—financial, operational, and commercial (Centre for Research on Energy and Clean Air, 2023)—sanctions initially restricted the provision of these Western services. In May 2022, the European Union (EU) agreed to ban the use of these Western services on all Russian seaborne exports. As global oil demand is inelastic and Russia was producing about 11% of the world’s oil, governments feared that this policy would have caused a spike in the world oil price.²

As an alternative, the US suggested a price cap. After a contentious debate about the cap level, the G7, the EU, and allies (hereafter, the West) agreed to impose a price cap of \$60 per barrel beginning in December 2022; many, including Ukrainian officials, continue to argue that a cap closer to Russia’s marginal cost of extraction would punish Russia more.³

The cap applies to any seaborne oil from Russia transported using Western services; oil transported without Western services is exempt from the cap. To take advantage of this exemption, Russia began to assemble a fleet (hereafter “the shadow fleet”) that uses non-Western services in order to sell oil at prices above the cap. In addition to cap sales and fleet sales, limited enforcement

¹Our calculation is based on direct military costs of the war estimated by Shatz and Reach (2023), federal budget information from International Monetary Fund (2021), and Russia’s GDP (World Bank, 2023).

²Western discussions about such a boycott coincided with a surge in the oil price, raising concerns about market imbalance and inflation. Following the surge in the Brent oil price, peaking at \$137 per barrel on March 7, 2022, US Secretary of State Blinken assured that sanctions would be implemented “while making sure that there is an appropriate supply of oil in world markets.”

³See, e.g., *Politico*, December 5, 2023, available at <https://www.politico.eu/article/russia-oil-price-cap-ukraine-war-centre-research-energy-clean-air/>.

of the cap also leads Russia to occasionally cheat by selling oil at prices above the cap while using Western services. Violators are subject to punishment, but only if an audit reveals that they knowingly engaged in cheating. Recently, Western countries have implemented new policies to deter such cheating and to reduce the size of the shadow fleet (Brooks & Harris, 2024). The goal of these policies is once again to reduce Russia’s ability to finance the war while avoiding a price spike.

The price cap on Russian oil is a new, untested economic sanction. It is currently a subject of active public discussion, with experts recommending potential adjustments⁴ and application to more countries,⁵ and policy-makers currently considering to tighten the price cap.⁶ Since a targeted cap is without historical precedent, prediction of market responses to this policy, including effects on the sanctioned country and on the world price of oil, must rely on economic theory supplemented by calibrated simulations. The novelty of the policy quickly piqued the interest of economists—see, e.g., Johnson, Rachel, and Wolfram (2023b) for a pioneering analysis, among other examples discussed in our literature review below. Yet, none of this burgeoning economics literature has taken into account the central feature of Russia’s policy response: Russia’s expansion over time of the shadow fleet and its increasing substitution of that fleet for tankers using Western services.

In this paper, we build a dynamic equilibrium model that accounts for the expansion of the Russian shadow fleet. We then calibrate its parameters to reproduce observed facts in data and use it to simulate the outcomes of various policies intended to reduce Russian oil profits.⁷ We say “intended” because we have discovered that some policies designed to shrink Russian profits do lower profits in the very short run but, in fact, raise future profits so much that the present value of the entire stream of profits increases.

In our dynamic model, Russia maximizes its discounted profits from oil exports taking as given

⁴See, e.g., Simon Johnson and Oleg Ustenko’s suggestion that the next US Administration should lower the price cap in *Newsweek*, December 4, 2024, available at <https://www.newsweek.com/trump-ukraine-plan-straightforward-nobel-prize-winner-1995598>. See also Spiro, Wachtmeister, and Gars’ (2024) comprehensive review of policy options.

⁵See, e.g., Simon Johnson and Catherine Wolfram’s call for imposing a price cap on Iranian oil exports in *The Washington Post*, October 18, 2023, available at <https://www.washingtonpost.com/opinions/2023/10/18/oil-iran-russia-war>, echoed by Jeffrey Sonnenfeld and Steven Tian who further suggest that this cap should be lower than \$60 per barrel imposed to Russia in *Yale Insights*, October 25, 2023, available at <https://insights.yale.edu/insights/to-prevent-wider-war-in-the-middle-east-choke-off-irans-oil-sales>.

⁶On G7 current discussions about lowering the cap or switching to a complete service ban, see, e.g., *Bloomberg*, December 19, 2024, available at <https://www.bloomberg.com/news/articles/2024-12-19/g-7-considering-options-to-harden-price-cap-on-russian-oil>.

⁷We focus on profits from oil exports both because they measure exports’ value added to the overall Russian economy and because the Russian government has increasingly oriented its oil tax system towards the taxation of profits since 2019 (Yermakov, 2024).

the world price path of oil. In each period, Russia can produce oil for export through two channels. It exports some oil at the cap using Western services and some at the world price minus a small discount using the shadow fleet. The world price path adjusts to clear markets. To isolate the effects of shadow fleet expansion, we assume that other sources of supply are constant⁸ and take no account of Russian oil being exhaustible.⁹

We assume provisionally that the cap is perfectly enforced. When we relax this assumption, there will be a third export channel—oil exported using Western services but sold at a producer price above the cap.

We assume the cap is binding and sufficiently tight that Russia always utilizes the shadow fleet. Initially, Russia has limited shadow-fleet capacity because, prior to the imposition of the cap, Russia relied on Western services. Given the limited initial capacity of the fleet, the marginal cost of additional production is below the \$60 cap, so Russia also finds it profitable to supplement its fleet sales with exports at the cap. Exports through the two channels are smaller than before the cap was imposed since the cap is binding. Consequently, the world price is higher than prior to the imposition of the cap. As fleet capacity expands, fleet exports replace cap exports leaving aggregate exports and the world price unchanged. However, cap sales eventually cease altogether and, since the fleet expansion continues, aggregate Russian exports begin to increase, driving down the world oil price.

We use our model to compare the effects of the service ban proposed by the European Union to the \$60 price cap ultimately adopted. The service ban can be viewed as equivalent to a price cap set so low (below \$34.1) that Russia would sell nothing at the cap. Russian profits under the ban would at first be considerably less than under the \$60 cap because of the massive reduction in exports (from approximately 5 mb/d to 2 mb/d). At the same time, the ban would impel Russia to expand its shadow fleet more rapidly than under the \$60 cap. In our simulations, after only one quarter, Russian profits are higher under the ban than under the cap, and profits remain higher

⁸We discuss how a positively-sloped non-Russian supply function affects the price elasticity of the residual demand facing Russia in Appendix B and show that this alternative assumption does not affect our findings that the cap and the service ban deliver similar present values of profits. However, were non-Russian supply sufficiently elastic, the ordering of policy would be reversed so that the service ban would harm Russia slightly more than the price cap.

⁹We also make this simplifying assumption for calibration purposes. In any case, while the standard model of exhaustible resources (Hotelling, 1931) is clearly appropriate for minerals that may be costly to extract, its applicability to oil is more controversial since, once an oil well is drilled, production from it does not cease immediately as the Hotelling model assumes but declines gradually over time.

through the 8th quarter. It turns out that the present value of Russian profits is smaller under the cap than under the ban. This difference is relatively small ($\approx 2\%$) so the main advantage of the price cap over the service ban is that, while both policies reduce Russia’s profits by more than 23% compared to no policy, the cap accomplishes this without creating a spike in the world price of oil.

The paradox is that Russia maximizes the present value of its profits when facing a \$60 cap but would earn a higher present value if instead it chose to act as if facing a service ban, which is also feasible. Such a paradox can occur under other assumptions about market structure—unless one makes the unrealistic assumption that Russia unilaterally controls the world price either as a monopolist or as a price-setting dominant firm.¹⁰ For example, a similar paradox occurs in Turner and Sappington (2024)’s Cournot duopoly model in which marginally lowering the cap below \$60 increases Russia’s profits.¹¹ The paradox, therefore, cannot be attributed to our price-taking assumption.

In addition to slightly lowering Russia’s discounted profits more than a ban, the price cap also hurts consumers significantly less. Sensitivity analyses (Appendix B) show that both results are robust to changes in our baseline assumptions about Russia’s discount rate and its marginal cost of fleet expansion.

We then investigate the impact of setting the cap at different levels. Our simulations indicate that the rate of expansion of the shadow fleet is sensitive to the cap level for caps between approximately \$34 and \$71 per barrel. In the interval between \$34 and \$69, lowering the cap *raises* the present value of Russian profits but only slightly whether we assume a yearly discount rate as low as 7.5% or as high as 30%. This implies that heated and divisive debates about the level of the cap are unwarranted. Lower caps punish Russia in the 1st quarter, when the size of the shadow fleet has not had time to adjust, but stimulate such a rapid expansion of this fleet that sales at the cap are soon eliminated. In general, the lower the cap, the larger the fleet size will be after the 1st quarter. A cap closer to the net price Russia receives on exports via the shadow fleet delays the fleet expansion and postpones the date when cap sales are replaced altogether. Curiously, a cap of

¹⁰In Appendix A, we analyze the unrealistic polar case in which Russia controls the world price. We find that the present value of the monopolist’s profits does not increase as the cap tightens, so that the paradox does not occur. However, this analysis shows that the profit increases seen in price-taking behavior are minimal, so any tightly binding cap delivers quantitatively similar profit losses.

¹¹In Turner and Sappington (2024) imposition of a \$60 price cap *raises* Russian profits compared to no policy. In addition, the imposition of the \$60 cap is predicted to increase exports sold at the cap using Western services. Neither result occurs in our model.

\$69.35 minimizes the present value of Russian profits.

Readers expecting reductions of the cap to lower Russia’s profits likely have in mind the profit consequences of decreasing a single price when all other prices facing a price-taking profit-maximizer remain constant. Hotelling’s lemma then implies that the firm’s profits must fall when the ceiling is reduced.¹² But with a second price endogenously determined in the free market, these familiar results do not apply. In this application, reductions in the cap cause sales at the ceiling to decrease, driving up the world price. When the ceiling is sufficiently low (below \$69.35 in our calibration), sales at the cap will be small relative to shadow fleet sales on the free market. A reduction in the cap in this circumstance must raise Russian profits since the increase in revenue from selling at an increased world price will exceed the decrease in revenue from cap sales at the reduced ceiling. At the end of Section 4, we illustrate graphically these opposing effects on profits and provide a necessary and sufficient condition for a tighter cap to benefit Russia. When we relax our baseline assumption of perfect enforcement, we find that this same condition also determines whether increased enforcement of the cap benefits Russia.

Finally, we investigate the effects of targeting the shadow fleet. We simulate an unanticipated reduction of the shadow fleet in the 12th quarter, which lowers the fleet size to where it had been in the 4th quarter (a reduction of approximately 35%). If this reduction in fleet size occurs while Western services are being used for some shipments, then there will be no increase in the world price since exports at the ceiling price will expand to offset the loss in sales at the net world price using the shadow fleet. In this case, the sanction must reduce the present value of Russia’s profits. However, if no Western services are being used immediately after the reduction in fleet capacity, then the world price will jump up and surprisingly, in our calibrated simulations, the present value of Russian profits increases.

Literature and contribution. Our analysis adds to the nascent literature that examines the price cap on Russian oil. An early and influential paper in this literature is due to Johnson et al. (2023b). They present a rich model in which Russia extracts oil to maximize an intertemporal objective that includes stochastic price variation and the need for stable revenues. Their assessment

¹²Hotelling’s lemma implies that the profit of a multiproduct firm must decline if the price of one product falls and the prices of its other products remain unchanged. It applies in our model only if there is no shadow fleet and so Russia can only sell at the ceiling price. See Section 4.4 for further details.

points to the usefulness of the price cap in limiting market power and stabilizing oil prices. It also highlights that the existence of a shadow fleet allowing Russia to evade the cap is critical to the effectiveness of the instrument. Yet, they leave the expansion of this fleet unexplored. Although our models are very different,¹³ we incorporate elements that address several issues they identify as important future steps, including the endogenous expansion of the Russian fleet.¹⁴

Other papers examine various aspects of the impact of the price cap. Assuming no evasion possibility, Wachtmeister, Gars, and Spiro (2023) conclude that Russia’s profits decrease as the cap is lowered. They further estimate that improvements in the Russian consumer surplus are relatively small. Johnson, Rachel, and Wolfram (2023a) examine the effect of price caps lower than \$60 in a static setting. Like Johnson, Rachel, and Wolfram (2023b), they assume a fixed shadow fleet capacity. Their model also implies that Russia’s profits decrease as the cap is lowered. They further suggest that Russia may stop cap exports at price caps below \$45 depending on the size of the shadow fleet. Both papers imply that a service ban would have lowered Russian profits more than a \$60 cap. Neither paper takes into account how lower caps hasten the expansion of the shadow fleet and how this affects Russian profits.

Building on Becko’s (2024) study of the effective structure of trade sanctions, Turner and Sappington (2024) present a static Cournot model in which Russia is a duopolist choosing export volumes sold under the cap and how much through its fleet, with fleet capacity being adjusted instantaneously. The introduction of the \$60 cap increases not only Russia’s fleet sales but also its exports using Western services. However, the latter prediction is at odds with observed behavior. Since the imposition of the \$60 cap also increases Russia’s oil profit, their analysis implies that not sanctioning would have been a better policy choice.¹⁵ In contrast, the imposition of a \$60 cap in our model reduces Russia’s sales using Western services and reduces the present value of Russian profits by 25%.

¹³A major difference is that we do not treat oil as an exhaustible resource. For simplicity, as well as for calibration purposes, we adopt a static model of oil supply. Although the rich structure of Johnson et al.’s (2023b) theoretical framework avoids the usual criticisms of Hotelling (1931) model, its adoption in our parsimonious setup would be problematic.

¹⁴As Johnson et al. (2023b) put it: “Several issues would be worth pursuing in the next stage. . . it would be useful to incorporate supply constraints or *convex costs of production* into our setting and think more carefully about the incentives to build up the shadow fleet, by making its *size endogenous* and determined optimally, perhaps subject to *adjustment costs*.” (Johnson, Rachel, & Wolfram, 2023b, p. 49) [our italics]

¹⁵In their analysis, revenue, profits, and sales using Western services are strictly increasing in the cap level for caps below \$56.35. A very low cap can reduce Russia’s profits by up to 19%.

Other papers in this literature have examined the cap following a different empirical approach. Babina et al. (2023) and Hilgenstock et al. (2023) use transaction-level data on Russian exports to provide evidence that Russia had to accept significantly lower producer prices in new markets. Kilian, Rapson, and Schipper (2024a) study the revenue loss due to observed discounts on Russian oil shipped through the shadow fleet. Their estimates—see also Wolfram’s (2024) subsequent comment and Kilian et al.’s (2024b) reply—suggest that evading the price cap through the fleet comes at various costs that fall on Russia, including augmented transportation costs through longer distances (e.g., to India), insurance premia (e.g., augmented risk of oil damages), and monopsony power of buyers of Russian oil transported using its shadow fleet. In our model, we consolidate these costs into an exogenous fixed cost per barrel shipped.

We are the first in this literature to take account of Russia’s expansion of its shadow fleet over time and to show, using a calibrated model, how the fleet expansion alters the effects of tightening the cap, increasing enforcement, and sanctioning the shadow fleet.

Beyond this literature on restricting the use of Western services, our analysis illustrates Mancur Olson’s (1962, 1963) explanation for why targeted economic sanctions often have a limited impact. Analysts often target inputs that an adversary was relying upon before hostilities commenced as essential, only to discover that they can be readily replaced by substitutes. A classic example is German ball bearings, which were assessed as essential to munitions production at the start of World War II. Although most German ball bearing factories were destroyed in massive Allied bombing campaigns, German munitions production continued unabated because, when the need arose, the Germans found substitutes. The rapid expansion of Russia’s shadow fleet shows that Western services proved not to be essential to Russia. Yet its substitution for Western services is costly and takes time. In our simulations, Russia’s losses when the price cap is imposed reflect the cost of this substitution. Moreover, our comparison between the price cap and a complete ban—more radical but less effective—highlights that various instruments differ in how fast they trigger targeted countries to adapt, with implications for the choice of effective sanctions.

We proceed as follows. In Section 2, we describe our model. In Section 3, we describe its calibration. In Section 4, we use our calibrated model to compare the effectiveness of a price cap vs. a service ban sanctions as well as the consequences of tightening a price cap under different enforcement levels; the section concludes by explaining theoretically why such tightening raises the

present value of Russia’s profits in our calibrated simulations. In Section 5, we use our calibrated model to illuminate the consequences of disqualifying (“targeting”) a segment of the shadow fleet while either maintaining or simultaneously changing the level of enforcement. Section 6 concludes the paper.

2 The Model

2.1 Assumptions

We assume Russia sells its oil to the global market. Some of this oil is sold at a price less than or equal to the price cap using Western services. We assume provisionally that the price cap is perfectly enforced.¹⁶

Russia sells the rest of its oil at a price higher than the cap using an aging shadow fleet. The shadow fleet has been assembled from a limited pool of old and poorly maintained tankers. The fleet is costly to operate since it relies on less efficient services, including non-conventional insurance, and is subject to an environmental risk premium because of the increased likelihood of oil spills. The shadow fleet also serves longer routes to access new markets. We capture these costs by assuming an exogenous cost per barrel shipped using the shadow fleet so that Russia receives for each barrel that its fleet exports the world price minus a constant discount.¹⁷

The world price is determined by supplies coming from Russia and from other countries. We assume that Russia’s time horizon is finite and that it takes the sequence of world prices as given. We assume neither the demand function nor non-Russian supply fluctuates over time. Therefore, the dynamics in our model are attributable entirely to the expansion of shadow fleet capacity.

Expanding the shadow fleet is costly for Russia, regardless of whether it owns the tankers or leases them under long-term contracts. Besides acquiring old tankers or contracting with their owners, investments required by the expansion of its shadow fleet include port improvements. We model the fleet as capital and its expansion as costly investment.

¹⁶We follow the literature in treating crude oil and its products as a homogeneous good that we refer to as oil. In February 2023, the cap on Russian crude oil was completed by similar price caps for oil products so as to avoid internal arbitrage. Accordingly, we refer to the set of price caps simply as *the* price cap on Russian oil.

¹⁷Although the monopsony power of importers of Russian oil delivered using the shadow fleet is often mentioned, we do not take it into account explicitly. Kilian, Rapson, and Schipper (2024a) suggest that it accounts for a declining part of the total cost of evading sanctions through the shadow fleet.

Finally, in line with the related literature, we rule out the possibility that Russia can store oil since its on-shore and off-shore storage capacity is limited (Johnson, Rachel, & Wolfram, 2023b).

2.2 Notation and Curvature of Functions

We adopt the following notation:

- K_t : the capacity of the Russian shadow fleet in period t .
- \bar{k} : the initial capacity of the Russian shadow fleet.
- I_t : Increase in the capacity of the shadow fleet between period t and $t + 1$.
- Q_t : Russian sales at or below the cap in period t .
- X_t : Russian sales in period t using the shadow fleet ($X_t \leq K_t$).
- R_t : the sum of Russian oil exports in period t through the two channels ($R_t = X_t + Q_t$).
- Z : Non-Russian sales in period t .
- $C(R_t)$: Russian total cost of producing crude to be exported ($C'(0) \geq 0, C'(R_t) > 0, C''(R_t) > 0$).
- $F(I_t)$: Cost of expanding the capacity of the shadow fleet during period t ($F'(0) = 0$; and $F'(I_t) > 0, F''(I_t) > 0$) for $I_t > 0$.
- $p(Z + R_t)$: world inverse demand for crude evaluated in period t ($p'(Z + R_t) < 0$).
- d : price discount per barrel on shadow fleet sales.
- \hat{p} : price cap.
- β : constant discount factor.
- R_c : the solution to $\hat{p} = C'(R)$.

2.3 Optimization Problem

Russia takes the world price sequence $\{p_t\}_{t=1}^T$ as given and maximizes the sum of its discounted profits with respect to $Q_t \geq 0, X_t \geq 0, I_t \geq 0$, and $K_t > 0$:

$$\sum_{t=1}^T \beta^{t-1} [Q_t \min(p_t, \hat{p}) + (p_t - d)X_t - C(X_t + Q_t) - F(I_t)]$$

subject to $K_1 = \bar{k} > 0$, $K_{t+1} = K_t + I_t$ and $X_t \leq K_t$ for $t = 1, \dots, T$. That is, in period t Russia exports Q_t barrels at no more than the cap and X_t barrels at the net world price, where the latter exports cannot exceed the inherited shadow-fleet capacity, K_t . Since $K_t = \bar{k} + \sum_{s=1}^{t-1} I_s$ for $t = 2, \dots, T$, we can simplify the problem by substituting out of K_t .

Assigning the multipliers $\{\alpha_t\}_{t=1}^{t=T}$ to the constraints on X_t , the Lagrangean is:

$$\mathcal{L} = \sum_{t=1}^T \beta^{t-1} \left\{ \min(p_t, \hat{p}) Q_t + [p_t - d] X_t - C(X_t + Q_t) - F(I_t) + \alpha_t \left[\bar{k} + \sum_{s=1}^{t-1} I_s - X_t \right] \right\}$$

The following conditions must hold with complementary slackness (abbreviated as c.s.)¹⁸:

$$Q_t \geq 0, \quad \min(p_t, \hat{p}) - C'(Q_t + X_t) \leq 0, \quad c.s. \quad (1)$$

$$X_t \geq 0, \quad (p_t - d) - C'(Q_t + X_t) - \alpha_t \leq 0, \quad c.s. \quad (2)$$

$$I_t \geq 0, \quad -F'(I_t) + \sum_{s=t+1}^T \beta^{s-t} \alpha_s \leq 0, \quad c.s. \quad (3)$$

$$\alpha_t \geq 0, \quad \left[\bar{k} + \sum_{s=1}^{t-1} I_s - X_t \right] \geq 0, \quad c.s. \quad (4)$$

To deduce the conditions holding in the competitive equilibrium, we assume the market clears in each period by replacing p_t by $p(Z + Q_t + X_t)$.¹⁹

¹⁸That is, $x \geq 0$ and $y \leq 0$ (or $y \geq 0$) with complementary slackness means that in addition to these two weak inequalities, at least one of them holds as an equality.

¹⁹An alternative way to derive the equilibrium conditions is to maximize the following "fictitious" payoff function: $\sum_{t=1}^T \beta^{t-1} (H(K_t) + \hat{p}Q_t - C(K_t + Q_t) - F(I_t))$ subject to $K_1 = \bar{k}$ and $K_{t+1} = K_t + I_t$,

$$\text{where } H(K_t) = \begin{cases} K_t[p(Z + R_c) - d] & \text{for } K_t < R_c \\ R_c[p(Z + R_c) - d] + \int_{R_c}^{K_t} [p(Z + u) - d] du & \text{for } K_t \geq R_c \end{cases} \quad (5)$$

with R_c defined as the solution to $C'(R) = \hat{p}$. $H(K_t)$ is continuously differentiable and concave. Since this concave optimization problem has a unique solution and the constraint qualification is satisfied, its complementary slackness conditions have a unique solution. Since these conditions must hold in any competitive equilibrium, a unique equilibrium exists. Stokey, Lucas, and Prescott (1989) emphasize that, with the exception of Becker (1985), few examples have been discovered where a dynamic optimization problem has the same first-order conditions as a competitive




Figure 1: Equilibrium depends on whether cap is non-binding ($\hat{p} > p_{LF}$), slightly binding ($\hat{p} \in [p^*, p_{LF}]$), or tightly binding ($\hat{p} \in [C'(0), p^*]$).

2.3.1 The Cap Set in One of Three Intervals

The equilibrium depends on the level of the cap (\hat{p}). The cap can be (1) non-binding, (2) slightly binding, or (3) tightly binding. The discussion below will establish that only if the cap is tightly binding will Russia use the shadow fleet. Since Russia relies heavily on this fleet, we discard the other two cases as unrealistic.

We use Figure 1 to discuss the three cases. In the figure, three curves are plotted against Russian aggregate oil exports in period t . The upward-sloping curve is Russia's marginal cost of production, $C'(R_t)$. The higher of the downward-sloping curves gives the world price as a function of Russian aggregate sales, $p(Z + R_t)$. The lower downward-sloping curve is this inverse demand curve shifted down by the exogenous constant discount, $p(Z + R_t) - d$. We refer to it as the “net price” curve.

The marginal cost curve intersects the inverse demand curve at the point (R_{LF}, p_{LF}) and the lower net price curve at the point (R^*, p^*) . Since the marginal cost curve is upward sloping, $p^* < p_{LF}$ and $R^* < R_{LF}$.

If the cap is set *above* p^* , it is either not binding ($\hat{p} > p_{LF}$) or only slightly binding, $\hat{p} \in (p^*, p_{LF})$.²⁰ Since the price in the two cases strictly exceeds what Russia would receive using the

equilibrium distorted by government intervention.

²⁰The “slightly binding” case disappears if $d = 0$. In that more conventional situation covered in the textbooks,

shadow fleet, Russia would never use the shadow fleet if the cap were set in either interval.

In reality, Russia has relied increasingly on the shadow fleet since the price cap was first imposed. We infer, therefore, that the cap is tightly binding ($\hat{p} < p^*$). In this case, the intersection of the cap with the marginal cost curve must lie strictly below the net price curve: $C'(R_t) < p(Z + R_t) - d$, where $R_t = X_t + Q_t$. Equation (2) then implies that $\alpha_t > 0$ and equation (4) implies that $[\bar{k} + \sum_{s=1}^{t-1} I_s - X_t] = 0$. Hence, $X_t = K_t$.

Since sales using the shadow fleet are more lucrative than sales at the cap, Russia will use the full capacity of the shadow fleet in every period. The initial capacity of the fleet (\bar{k}) is assumed to be small since Russia was relying on Western services prior to the sanctions. If the marginal cost of expanding production *beyond* the initial capacity is smaller than the additional revenue from selling at the cap ($C'(\bar{k}) < \hat{p}$), additional sales using Western services will occur to the point where the marginal cost of producing for export through the two channels equals the cap (equation (1)). Russia will export K_t using the fleet and $R_c - K_t$ using Western services where R_c solves $\hat{p} = C'(R)$. The lower the cap, the smaller R_c and the higher $p(Z + R_c)$.

Hence, through the two channels, Russia sells in aggregate R_c barrels. As sales using the shadow fleet expand, sales using Western services decline by the same amount so that aggregate Russian sales continue to be R_c and the world price continues to be $p(Z + R_c)$. Once the capacity of the shadow fleet reaches R_c , however, sales at the cap cease ($Q_t = 0$); further expansion of the shadow fleet capacity reduces the world price monotonically.

Let R^* be defined as the unique solution to $p(Z + R) - d = C'(R)$ and $p^* = p(Z + R^*)$. The shadow fleet capacity will never exceed R^* . If T is sufficiently large, fleet capacity will converge to R^* from below and the net price will converge to p^* from above. Neither R^* nor p^* depends on the size of the cap.

3 Model Parametrization and Calibration

To make the model numerically tractable, we further specify parameter values and functional forms. This section outlines the specifications and rationale for our choices. Key calibrated parameters are presented in Table 1. In our model, time periods are quarters. Our baseline simulations adopt the cap either binds or does not bind.

a deterministic termination period of $T = 80$ quarters (20 years) and an annual discount rate of 15% ($\beta \approx 0.96$). We choose this discount rate to approximately reflect Russia’s cost of capital, as indicated by the yield rate on zero-coupon Russian government bonds (Central Bank of Russia, 2024). Nevertheless, we examine alternative values for these parameters in sensitivity analyses and find that our qualitative results are robust to a wide range of discount rates (see Appendix B).

Parameter	Variable	Value	Source
Sanction termination	T	80 quarters	Modeling assumption
Discount factor	β	0.960	Modeling assumption
Non-Russian supply	Z	92.0 mb/d	IEA, Gars, Spiro, and Wachtmeister (2022)
Initial capacity	\bar{k}	2.0 mb/d	IEA, modeling assumptions
Initial supply to non-coalition	Q_0	5.4 mb/d	IEA, modeling assumptions
Pre-invasion crude price	p_0	\$80	IEA (approximated)
Shadow fleet discount	d	\$15	Neste/Thomson Reuters
Price elasticity of oil demand	ϵ_D	-0.125	Gars et al. (2022)
Marginal cost intercept	c_0	\$17	Rystad, Gars et al. (2022)
Marginal cost slope	γ	0.095	Calibrated to match IEA data
Marginal cost of investment slope	ϕ	4.102	Calibrated to match IEA data

Table 1: Parameters used in baseline simulations.

We set values for initial export volumes based on IEA reports for the year 2021—before the invasion of Ukraine (International Energy Agency, 2024). We aggregate four major importers of Russian oil post-invasion—China, India, Turkey, and the Middle East—and set the initial evasion capacity equal to the daily average volume of Russian exports to these regions in 2021: $\bar{k} = 2.0$ million barrels per day (mb/d). As a result, initial exports to all other regions (Q_0) are equal to 5.4 mb/d. As in Gars et al. (2022), we use IEA data to determine the non-Russian global supply of oil, with $Z = 92.0$ mb/d assumed constant throughout. Moreover, we adopt a pre-invasion reference crude price of $p_0 = \$80$ per barrel. In order to gauge the impact of sanctions, we will often compare our simulations to a Business-as-Usual (BAU) case before the imposition of the import embargo and other sanctions. We define BAU as the repetition of $t = 0$ prices and quantities during all simulated quarters.

Our baseline simulations are based on a fixed discount level of $d = \$15$ per barrel for exports using the shadow fleet. The baseline d is the approximate value at which the Brent-Urals differential stabilized one year after the start of sanctions. This value is in line with Kilian, Rapson, and Schipper (2024a), who find that increased shipping distances lead to a discount of \$12–\$15 per barrel.

Three functions need to be specified in the simulations: global inverse demand for oil (p), production costs (C), and investment costs (F). Our specifications favor parsimonious functional forms that capture stylized economic facts and align with related papers in the literature (Gars, Spiro, & Wachtmeister, 2022; Wachtmeister, Gars, & Spiro, 2023; Johnson, Rachel, & Wolfram, 2023a, 2023b).

We model global inverse demand for oil as an isoelastic form, given by

$$p(Z + Q_t + K_t) = p_0 \left(\frac{Z + Q_t + K_t}{Z + Q_0 + K_0} \right)^{1/\epsilon_D} \quad (6)$$

where ϵ_D is the price elasticity of demand. Our baseline case in the simulations follows Gars, Spiro, and Wachtmeister (2022) and adopts a price elasticity $\epsilon_D = -0.125$; this value is supported by recent empirical estimates (see Kilian (2022) for a review), but we consider alternative choices in the sensitivity analyses in Appendix B. While the world demand curve is assumed to be inelastic, the residual demand curve facing Russia is elastic since less than 12.5% of world demand is satisfied by Russia. Indeed, since there is less demand at higher prices and non-Russian sources are always assumed to satisfy Z of it, the magnitude of the elasticity of the residual demand rises with price and residual demand is zero at prices above approximately \$150 ($p(Z) \approx \150).

Production costs are given by $C(R_t) = c_0 R_t + \gamma \frac{R_t^2}{2}$, where c_0 and γ are, respectively, the intercept and slope of the linear marginal cost of production curve. Following estimates in Gars, Spiro, and Wachtmeister (2022) based on Rystad Energy data, we adopt $c_0 = \$17$ as the vertical intercept of the marginal cost. The slope parameter γ is calibrated using first-order optimality conditions in the pre-invasion period. Our model assumes that Russia is a price taker, in which case the condition is given by $p_0 = C'(R_0)$, thus yielding $\gamma = \frac{p_0 - c_0}{R_0}$. Using the initial price and quantities outlined above results in a calibrated $\gamma \approx 0.095$.

Sales at the cap begin in the 1st quarter and cease on the quarter before R_c is first surpassed. For the cap of \$60, $R_c \approx 5.0$ mb/d and the net price plateau from the initial quarter remains $p^* = 81.87$. A lower cap would result in a smaller value for R_c and a higher net price plateau.

The marginal-cost curve intersects the net price curve at $R^* = 6.4$ mb/d. The net-of-discount price is then $p^* = 71.63$.

The investment cost function is also quadratic, given by $F(I_t) = \phi \frac{I_t^2}{2}$, where ϕ is the slope of the

marginal cost of investment. Hence $F'(0) = 0$. Given this assumption, $I_t > 0$ until the penultimate quarter ($T - 1$), no matter the size of T . We calibrate ϕ using a numerical root-finding method.²¹ This method searches for the parameter value in which the solution of the model gives a shadow fleet capacity at the end of the second year that matches the IEA data ($K_9 = 5.2$ mb/d). In our baseline simulations, this approach results in a calibrated value of $\phi \approx 4.1$.

4 Effects of Reducing the Payoff per Barrel on Sales Using Western Services

This section consists of four subsections. The first three show that different policies widely believed to harm Russia would in fact raise the present value of its profits back towards their BAU level. In the first subsection, we compare the present value of Russian profits when a \$60 cap is imposed and the profits when all sales using Western services are banned (as initially advocated by the EU).²² Surprisingly, in our simulations the \$60 cap harms Russia more than the service ban. Second, we consider the policy of lowering the cap, as many supporters of Ukraine advocate. In our simulations, lowering the cap below \$60 harms Russia less in terms of the present value of Russian profits. Third, we relax our assumption of complete enforcement. In our simulations, tighter enforcement lowers the expected payoff per barrel and, like a tighter ceiling when enforcement is perfect, hurts Russia less. The fourth subsection explains the source of these surprising results.

4.1 \$60 Cap vs. Service Ban

Instead of the outright ban on the use of Western services proposed by the EU, the US proposed allowing the use of such services but only if the oil transported was sold at a producer price not exceeding the price cap; under G-7 rules, oil sold at a higher price would have to be carried by the shadow fleet. To simulate the service ban using our model, we set the price cap so low ($\hat{p} \leq 34.1$) that Russia would choose not to use Western services and instead rely on its shadow fleet.

To gain intuition about the consequences of the two policies, suppose the initial fleet capacity

²¹Appendix B discusses alternative specifications of F , performs a sensitivity analysis on ϕ , and shows that different values do not change our qualitative findings.

²²We simulate a service ban by assuming a ceiling so low that Russia would choose to sell nothing at that price. Since $C'(\bar{k}) = 34.1$, any ceiling below \$34.1 would reward Russia less per barrel sold at the ceiling than the marginal cost of producing it.

were zero and that Russia had been completely reliant on Western services. Then exports would have been zero under the service ban but R_c under the cap. The removal of Russian oil from the market would have driven the world price to approximately \$150 ($p(Z) \approx 150$) but since Russia would have had no way to export, its net revenue (revenue net of production costs) would have been zero and, given its investment costs, its profit (net revenue *minus* investment costs) would have been negative. Clearly, the immediate effect of imposing a service ban would have been a loss in Russian profits compared to a price cap.

However, since Russia would have been unable to export under a service ban without expanding its shadow fleet, such expansion would have been more rapid than under a price cap. Eventually, Russian exports under a service ban would catch up with the exports under a price cap. In that situation, the world price under the service ban would have fallen to the constant price under the cap, and the production cost under the service ban would have risen to the constant level under the cap. However, since part of the exports under the cap sells at the cap price instead of the net world price, net revenue would be higher under the service ban.

We report our simulations of the two policies in Figure 2. In our calibration, \bar{k} is 2 mb/d (not zero, as in the foregoing thought experiment). Notice that Russian quarterly profits (in the bottom left panel) are initially smaller under the service ban (5.54 \$B/quarter under the ban and 8.86 \$B/quarter under the cap) but by the 2nd quarter, Russian profits are higher (11.72 \$B/quarter under the ban and 10.56 \$B/quarter under the cap) because the service ban stimulates a more rapid expansion of the shadow fleet (bottom right panel). Hence, it only takes one quarter for this profit reversal to occur.

The bottom left panel of Figure 2 also reports production profits in each quarter if neither of the policies was imposed and there was no import embargo (BAU). In that case, p_{LF} would prevail in every quarter, and Russian profits would be steady and much larger than under either of the proposed sanctions. The bottom right panel of Figure 2 summarizes the present values of profits under the two policies (and no policy). While the two policies harm consumers compared to no policy, the cap harms consumers less since it results in a lower price of $p(Z + R_c)$ for nearly eight quarters, beginning when the sanction is imposed. These results are robust to changes in the discount rate and the marginal cost of fleet expansion; however, a substantially higher price elasticity of oil demand facing Russia would lead to the present value of profits under the service

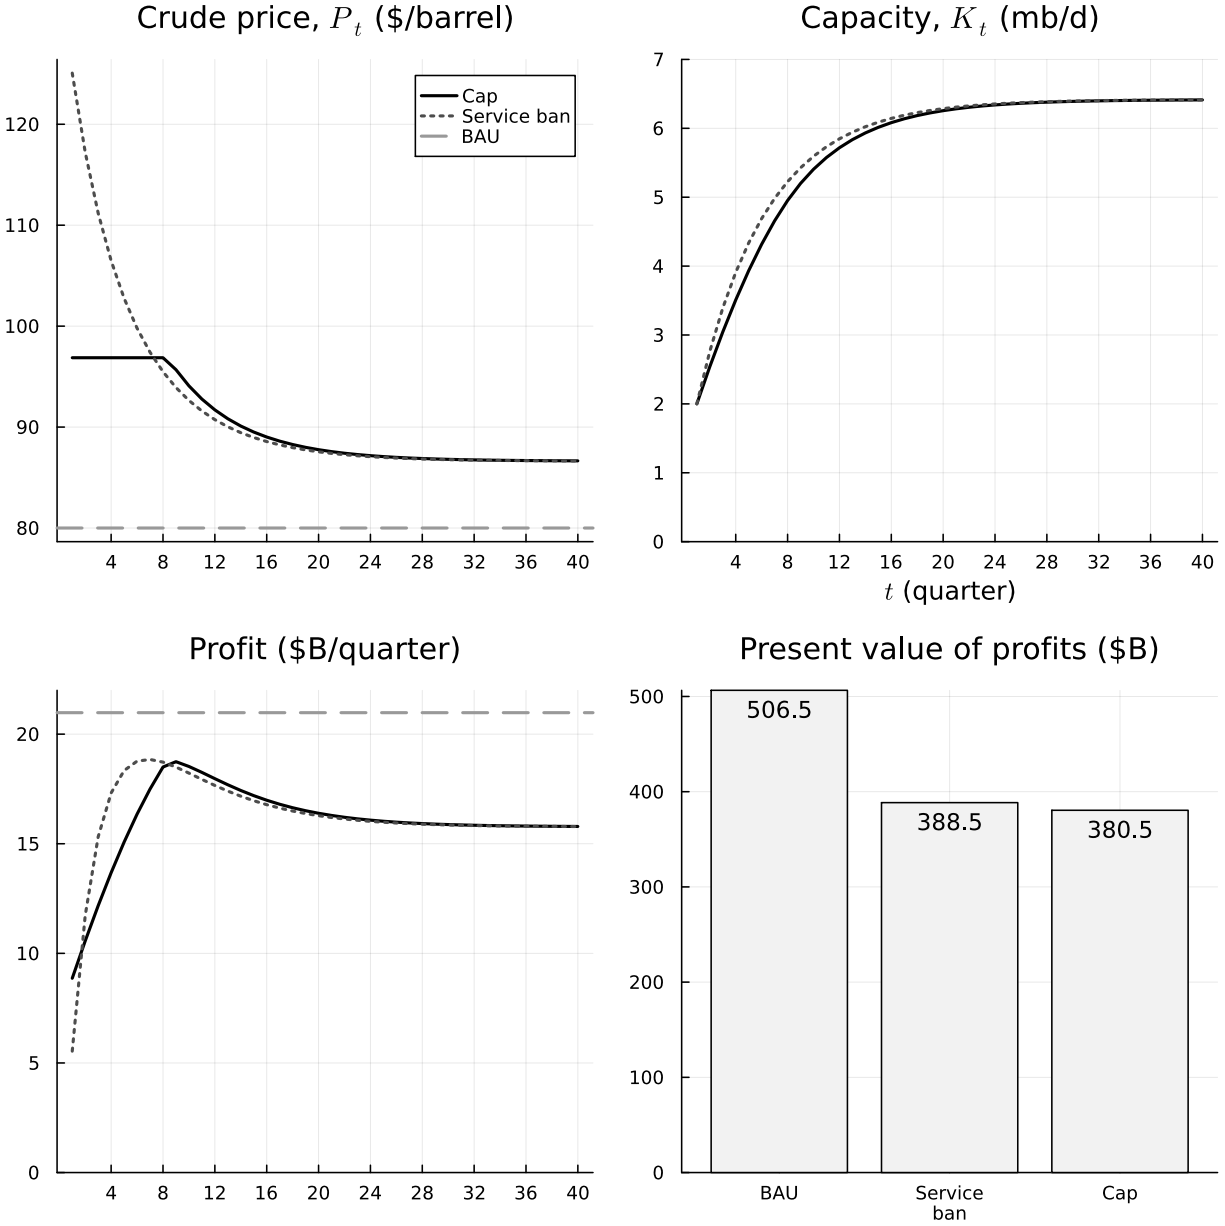


Figure 2: A comparison of prices, capacity, and profits under a price cap sanction (solid lines) and a service ban (dotted lines) The line graphs display only the first 40 quarters of the 80-quarter simulation.

ban being slightly lower than under the cap (see Appendix B for details).

4.2 Lowering the Cap

Our simulations confirm a more general result: For caps lower than \$60, the lower the cap, the higher the fleet capacity after the 1st quarter and the less it hurts Russia in terms of the present

value of Russia's stream of profits. We present these results in Figure 3.

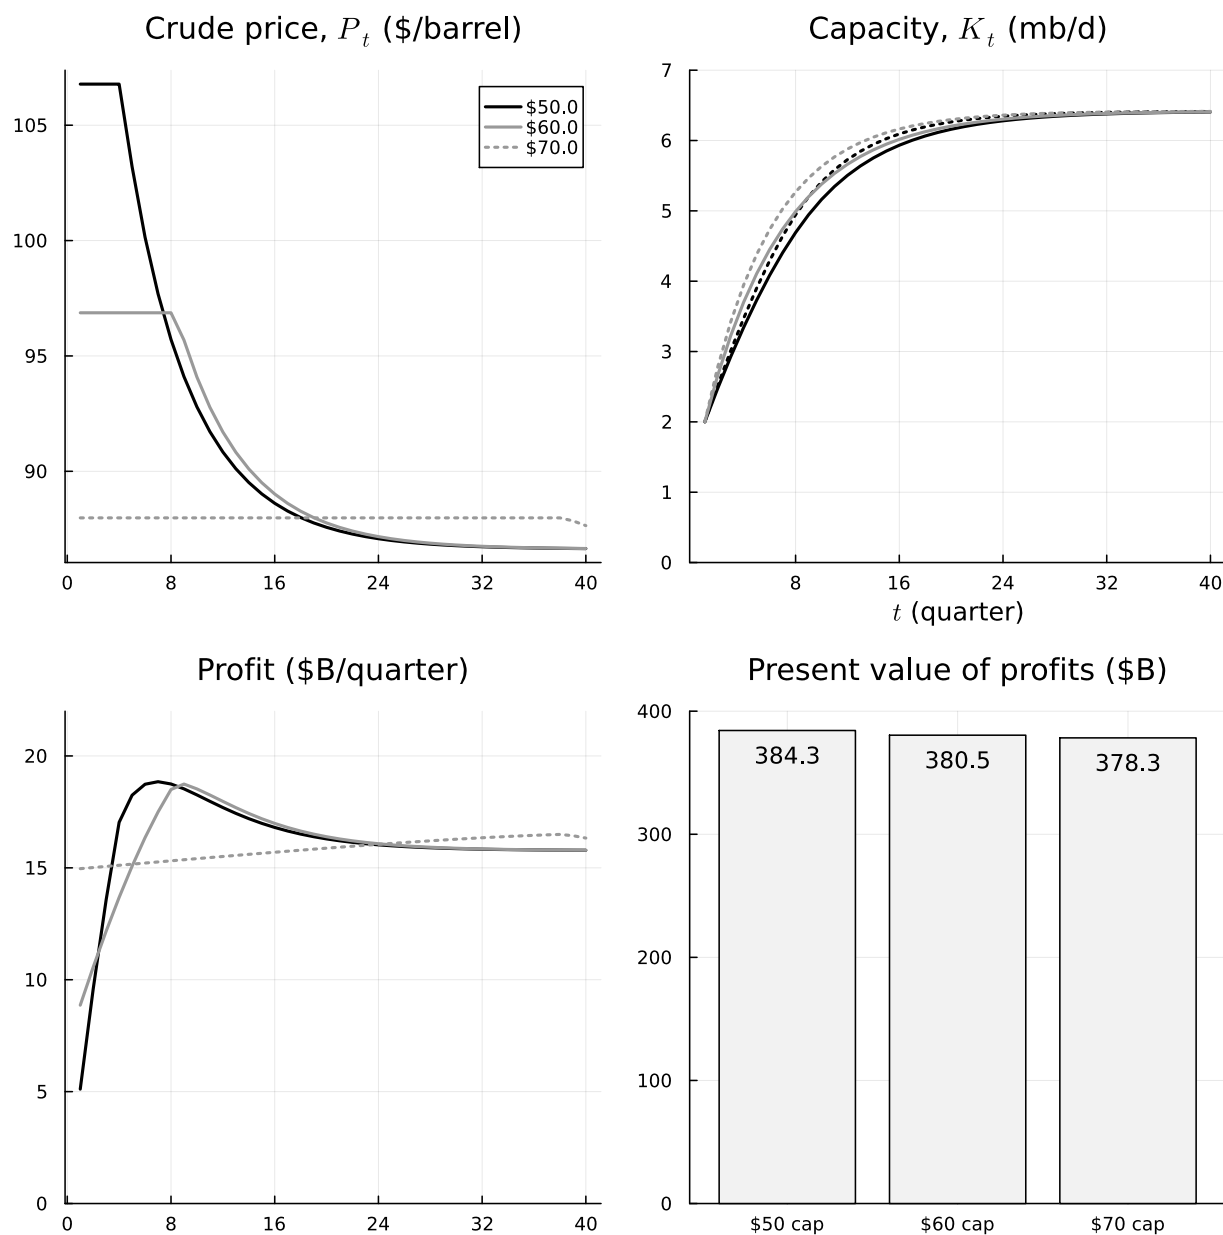


Figure 3: A comparison of prices, capacity, and profits under different cap levels. The line graphs display only the first 40 quarters of the 80-quarter simulation.

Hence, a cap lower than \$60 would hurt Russia less than a \$60 cap. Indeed, as our simulations with a \$70 cap suggest, a slightly higher cap would punish Russia more.²³ In general, a higher cap retards the time when Russia completely abandons its sales at the cap in favor of fleet sales, and

²³The present value of Russian oil profits is minimized at a \$69.35 cap. Although the present value is not monotonic in the cap level, the fleet size after the 1st quarter is monotonic in the cap level.

this advances the West’s goal of reducing Russia’s ability to finance the war.

4.3 Tightening Enforcement of the Cap

In Section 2, we assumed Russia never uses Western services when selling oil above the cap. We relax that assumption now. Suppose Russia exports Y_t barrels in period t on tankers that, although using Western services, carry oil costing more than the cap. Assume the enforcement authority can observe the price buyers pay, but determining insurance information, which may require detecting forgeries, requires a detailed audit. Some enforcement actions can be based solely on price information. The enforcer would never question cargoes sold in compliance with the cap and would always question shipments selling for a higher price other than $p_t - d$, the price charged for shadow fleet cargoes.

In response, Russia would price the Y_t barrels violating the agreement at $p_t - d$, rendering them indistinguishable—without an audit—from the X_t barrels shipped using the shadow fleet. The purchasers of these Y_t barrels also benefit from this arrangement since they can often escape punishment by insisting that they had thought the oil was being shipped on tankers using non-Western insurance.

Let $a \in [0, 1]$ denote the probability the enforcer audits a tanker carrying oil sold for $p_t - d$ per barrel. Whenever a tanker in the shadow fleet happens to be audited, it will never be penalized since it is in compliance with the agreement. However, whenever a tanker using Western insurance but carrying oil priced at $p_t - d$ is audited, it will be punished. We assume the penalty for violating the cap is (1) a monetary fine per barrel of τ^{24} and (2) the requirement that the oil be sold instead at the cap price (\hat{p}).

Since decisions about capacity expansion are analogous to those discussed in Section 2, we focus here mainly on the choice of Q_t , X_t , and Y_t . We denote Russian aggregate exports in period t as R_t , which now includes Y_t (assumed throughout Section 2 to be zero). We continue to assume the cap is set low enough ($\hat{p} < p^*$) that the full capacity of the shadow fleet is utilized ($X_t = K_t$) in every period.

²⁴In practice, it is up to the country conducting the audit to determine the form of the punishment. Some use fines, others use another sanction.




Figure 4: Audit probabilities lower than \hat{a} make cheating strictly more profitable than compliance with the cap; as the audit probability is lowered further, the expected profit increases and more cheating occurs.

In period t , Russia maximizes

$$\hat{p}Q_t + (p_t - d)K_t + [a(\hat{p} - \tau) + (1 - a)(p_t - d)]Y_t - C'(Q_t + Y_t + K_t). \quad (7)$$

Therefore, in equilibrium, the following four conditions define Q_t, X_t, Y_t , and R_t :

$$R_t = Q_t + Y_t + X_t, \quad (8)$$

$$X_t = K_t, \quad (9)$$

$$Q_t \geq 0, \quad \hat{p} - C'(R_t) \leq 0, \quad c.s., \quad (10)$$

$$Y_t \geq 0, \quad [a(\hat{p} - \tau) + (1 - a)(p(Z + R_t) - d)] - C'(R_t) \leq 0, \quad c.s.. \quad (11)$$

Figure 4 is useful in understanding these conditions. If $Q_t > 0$, condition (10) implies $R_t = \hat{R}$, the solution to $\hat{p} = C'(R_t)$. If $Y_t > 0$, condition (11) implies $R_t = R(a)$, the solution to $[a(\hat{p} - \tau) + (1 - a)(p(Z + R_t) - d)] = C'(R_t)$. If $R(a) < \hat{R}$, $Y_t = 0$ and $Q_t > 0$. This will occur, for example, if $a = 1$. Indeed, as Figure 4 illustrates, it will occur if $a > \hat{a}$, where \hat{a} solves $\hat{p} = a(\hat{p} - \tau) + (1 - a)p(Z + \hat{R})$. That is, if an audit is sufficiently likely, there will be no cheating ($Y_t = 0$). This corresponds to what we analyzed in Section 2.

Suppose, however, $a < \hat{a}$. Then $R(a) > \hat{R}$, $Y_t > 0$ and $Q_t = 0$. That is, if an audit is sufficiently unlikely, Russia will violate the cap. In this case, Y_t solves $a(\hat{p} - \tau) + (1 - a)p(Z + Y_t + K_t) =$

$C'(Y_t + K_t)$. Reductions in the audit probability raise the level of cheating ($\frac{dY_t}{da} < 0$) and Russia's expected profits.

If the audit probability is low enough to induce cheating ($a < \hat{a}$), total exports remain constant during the first phase; as a result, the world price also remains constant. The expansion of the shadow fleet leads to increased exports using non-Western services and an equal reduction in exports shipped in violation of the cap. When such cheating ceases altogether, the first phase ends. Thereafter, the shadow fleet continues to expand, driving down the world price. In this second phase, Russia relies entirely on its shadow fleet for its oil exports.

Figure 5 describes the dynamic consequences of alternative audit probabilities. The dashed line in each panel indicates what happens if the audit probability is small ($a = .2$). The black line in each panel represents the case where enforcement is perfect ($a = 1.0$, a reprise of Figure 2). The dotted line in each panel represents the case where the audit probability is at an intermediate level ($a = .5$) but still low enough to induce cheating. A lower audit probability results in a higher expected revenue per barrel, higher total exports during the first phase, a lower world price during that phase, and a longer duration of that phase. When the audit probability is nearly zero, cheating is almost as lucrative as exporting using the shadow fleet. Since the shadow fleet is capacity-constrained and expanding its capacity is costly, Russia would cheat initially and expand the shadow fleet very slowly.

In the short run, more stringent enforcement lowers Russian expected profits. But once again, this induces a faster expansion of the shadow fleet. As the top right panel of Figure 5 depicts, the more stringent the enforcement, the higher the fleet capacity after the 1st quarter.²⁵ The expected profit ranking eventually reverses, and in terms of the present value of expected profits, the more stringent the enforcement, the higher the sum of discounted Russian profits.

To illustrate, compare enforcement so stringent that no cheating occurs with enforcement so lax that only 20% of the tankers charging prices above the cap are audited. As Figure 5 reflects, Russian expected profits are initially much lower if enforcement is stringent. But stringent enforcement induces a rapid expansion of the shadow fleet. As a result, a profit reversal occurs. Between the 5th quarter and the 19th quarter (more than three years), Russian profits are actually higher if

²⁵While it may be difficult to see in the figures for large t that the fleet capacities never cross, a continuous-time argument can establish this point: If they ever crossed, they would have the same K_t at the same date and the same dynamics. Hence, they would have to stick together afterwards.

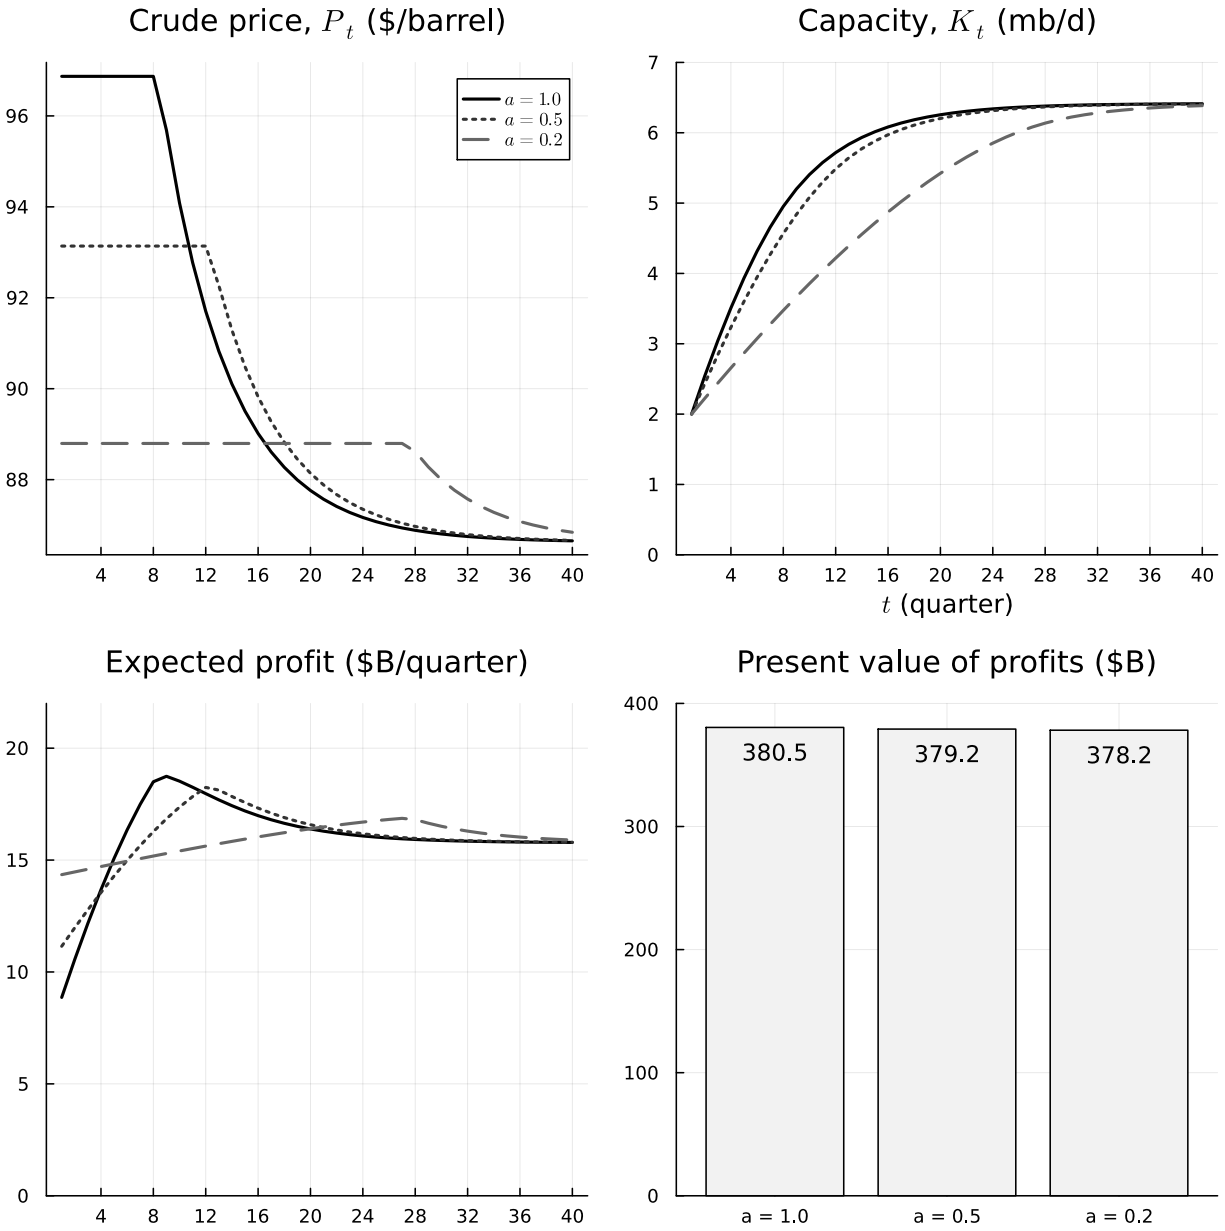


Figure 5: A comparison of outcomes under different levels of enforcement of the \$60 price cap sanction. Parameter a is the probability of an audit. All scenarios assume a price cap sanction and a monetary fee of $\tau = 10$ dollars/barrel. The line graphs display only the first 40 quarters of the 80-quarter simulation.

enforcement is stringent than if it is lax. So, surprisingly, lax enforcement hurts Russia more in terms of the present value of its expected profits.

4.4 Explaining the Surprising Results

In this subsection, we first deduce a necessary and sufficient condition for tightening the ceiling to raise static production profits in the 1st quarter; we then explain why this condition also implies that tightening the ceiling hurts Russia less in terms of the present value of Russian profits over the entire horizon.

When the cap is tightened, 1st-quarter sales at the ceiling decline and so does revenue from these sales; but with less oil on the world market, the world price increases, raising the revenue on sales from the capacity-constrained shadow fleet. In Figure 6, we depict the 1st-quarter gain as rectangular area A and the loss as trapezoidal area B . In our calibrated model, area A is larger.

A marginal reduction in the price ceiling will raise Russian profits in the 1st quarter if and only if:²⁶

$$\frac{Q_1}{\bar{k}} < -\frac{P'(Q_1 + \bar{k} + Z)}{\gamma}. \quad (12)$$

As the Figure and inequality (12) establish, the surprising result depends on the parameters in our calibration. For example, if the inherited fleet capacity (\bar{k}) were sufficiently small or the slope of the marginal cost of production (γ) were sufficiently large, the surprising result reviewed in the previous three subsections would not occur.

Suppose inequality (12) holds in the 1st quarter for some ceiling (\hat{p}). Then it must hold for lower ceilings as well (as long as $Q > 0$) since when the ceiling is lowered, Q_1 declines and $-P'$ strictly rises if the inverse demand curve is strictly convex and is unchanged if the inverse demand curve is linear. So the strict inequality (12) will continue to hold for any lower price ceiling.

Determining how a policy change affects production profits in the 1st quarter permits us to say what happens to the present value of profits (including F) over the entire horizon. In equilibrium, the present value of Russian profits depends on \hat{p} and ϕ , the slope of the marginal cost of expanding capacity ($F''(I) = \phi$).

Consider the present value of profit as a function of ϕ for two distinct values of \hat{p} .²⁷ Suppose

²⁶Since $\hat{p} = C'(Q_1 + \bar{k})$, $\frac{dQ_1}{d\hat{p}} = \frac{1}{C''}$. In equilibrium, 1st-quarter profits are: $\pi_1 = \hat{p}Q_1 + \bar{k}[P(Q_1 + \bar{k} + Z) - d] - C(Q + \bar{k})$. Differentiating, we conclude: $\frac{d\pi_1}{d\hat{p}} = [\hat{p} - C'(Q_1 + \bar{k}) + \bar{k}P'(Q_1 + \bar{k} + Z)]\frac{dQ_1}{d\hat{p}} + Q_1$. Using the first-order condition, the first two terms in the first factor on the right are zero. Using the comparative statics, $\frac{dQ_1}{d\hat{p}} = 1/C''$. Since we assumed $C(R) = .5\gamma R^2$, $C'' = \gamma$. $\frac{d\pi_1}{d\hat{p}} = Q + \bar{k}\frac{P'(Q_1 + \bar{k} + Z)}{\gamma}$. Hotelling's lemma arises as a special case if $\bar{k} = 0$. In that case, the rectangle in Figure 6 has no width and hence zero area. Inequality (12) follows by setting $\frac{d\pi_1}{d\hat{p}} < 0$.

²⁷We are indebted to Gérard Gaudet for suggesting that we analyze the extremes of profit as a function of ϕ for




Figure 6: Tightening the ceiling from \hat{P}_H to \hat{P}_L causes 1st-quarter sales using Western services to contract from Q_H to Q_L driving up the world price by $P_L^W - P_H^W > 0$. Hence, production profits on fleet sales using the inherited capacity increase by area A but profits on sales using Western services fall by area B . In our calibrated simulations, $A - B > 0$ although the difference is small.

$\phi = 0$. Whether or not we include F will not matter since $F = 0$. Suppose (12) holds and hence the tighter of the two caps results in strictly higher production profits in the 1st quarter. In the dynamic simulation, there would be full expansion of fleet capacity to R^* in the 2nd quarter since $\phi = 0$ and nothing would be sold at the cap thereafter. Since the two policies generate identical profits after the 1st quarter discounted at the same β and the tighter cap generates larger profits in the 1st quarter, it must generate larger present value of profits over the entire horizon.

Now consider ϕ so high that $I_t = 0$ for all t .²⁸ Once again $F = 0$ —this time because $I = 0$. Russia would in equilibrium simply repeat what it did in the 1st quarter in every subsequent quarter. As a result, the tighter ceiling would again generate a strictly higher present value of profits over the entire horizon. Indeed, given the repetition of the 1st-quarter advantage over the $T - 1$ quarters, the difference would be magnified by $\frac{1 - \beta^{T+1}}{1 - \beta} > 1$. Hence if ϕ is at either extreme, the policy generating the larger production profit in the first period generates a larger present value of profit (inclusive of F) over the whole horizon.

It can be shown analytically that the present value of profits is a continuous function of ϕ .

the cap and the service ban.

²⁸Since $F = .5\phi I^2$, there is no ϕ such that $I = 0$ although as $\phi \rightarrow \infty$, $F(I_t) \rightarrow 0$ for all t .

While it is in principle possible for the function with the tighter ceiling to cross the function with the looser ceiling, the two curves would have to cross an *even* number of times for (1) the tighter cap to be less profitable for Russia at some intermediate ϕ and yet (2) the present value of profits under the tighter cap to be higher at both extreme values of ϕ .

We have verified that in our calibrated simulations the two functions *never* cross (see Appendix B). This exercise also confirms that while our quantitative results are doubtless sensitive to ϕ , the surprising qualitative results in the previous three subsections do not depend on this parameter.

5 Targeting the shadow fleet

5.1 Understanding the effects of targeting

Policies have recently been considered to reduce the size of the shadow fleet. These include denying certain vessels access to ports or waterways as of time t^* or Ukraine's destruction of Russian tankers at t^* . Such targeting reduces shadow fleet capacity at $t^* + 1$ to a level K_{t^*+1} determined by the policy. Sometimes targeting is paired with changes in the enforcement policy. The goal is again to reduce Russia's present value of profits without causing a large spike in the world price of oil. We consider such policies in this section. We begin with targeting policies where enforcement remains unchanged.

It is useful to divide such targeting cases into three categories: (i) when Western services are being utilized at the time targeting is implemented; (ii) when only the shadow fleet is utilized after targeting is implemented; and (iii) when only the shadow fleet is utilized before targeting is implemented but tankers using Western services are added immediately after targeting is implemented.

As shown below, targeting in case (i) always reduces the present value of Russian profits. To show this, we first explain why targeting reduces production profits in the quarter when targeting is implemented and then show why this implies that the present value of profits must fall as well.

Whenever targeting is implemented during the phase when the world price is constant (case (i)), there would be no jump at all in the world price; the downward jump in exports using the shadow fleet would be matched by an offsetting upward jump in exports using Western services. This offsetting upward jump in the use of Western services could reflect either an expansion of

authorized sales at the cap or, alternatively, an expansion of unauthorized sales (cheating).²⁹ In either case, aggregate Russian exports would not change in response to targeting and neither would the world price. In that circumstance, production profits in the quarter when targeting occurs must fall since costs do not change and some output previously sold at a higher price is now sold at a lower one.

We can extend this insight about production profits in the 1st quarter of targeting to discounted profits over the remainder of the time horizon. If the cost of adjustment is zero ($\phi = 0$), then targeting reduces Russia's profits in the quarter when the policy is implemented; but then, regardless of whether there was targeting or not, Russia would fully adjust its shadow fleet capacity to R^* in the next quarter and per-quarter profits would be constant until the end of the horizon with or without targeting. Hence, targeting would lower the present value of profits because of its effect on profits in the quarter in which targeting was implemented. If, on the other hand, the cost of adjustment (ϕ) was so high that there was no adjustment ($I_t \approx 0$), then Russia would be harmed not only in the quarter when targeting was implemented but in every subsequent quarter since there would be no adjustment of shadow fleet capacity. So targeting would reduce the present value of Russia's profits from that quarter onward by a magnified amount relative to the case with $\phi = 0$. Unless the profit as a function of ϕ under the two policies (targeting and no targeting) crossed an even number of times (no crossing occurs in our simulations), targeting when there is no initial effect on the world price must harm the present value of Russia's profits regardless of ϕ .

Suppose instead the shadow fleet alone is used after (and hence immediately before) targeting (case (ii)). In that case, targeting will cause the fleet capacity to jump down and the world price to jump up. In principle, the effect on production profits when targeting occurs is ambiguous and so, by extension, is the effect on the present value of Russian profits. In that 1st quarter, Russia gains because the part of its exports that it continues to sell fetch a higher price but loses because it no longer sells the remaining output which had been earning a production profit. Either effect could dominate.

If losses dominate in the quarter of implementing targeting, then by the usual argument, tar-

²⁹Robin Brooks and Ben Harris (2024) point out that the sanctioning of 14 Sovcomflot tankers on February 23, 2024 eliminated their use in exporting Russian oil but generated no price spike. The authors speculate that no spike occurred because an expansion of exports using Western services offset the contraction of exports due to targeting without distinguishing between the authorized use of such services and cheating.

getting must reduce the present value of Russian profits. Indeed, even if targeting was large enough that Russia *resumed* use of Western services (case (iii)), production profits in the quarter when targeting occurs would have to fall since the entire capacity contraction can be broken into two steps. In the first, the capacity is reduced as far as possible without triggering the use of Western services; in the second, fleet capacity is reduced the rest of the way to complete the targeting policy. Russian production profits would decline on each step and hence would decline overall. So, by our usual argument, the present value of Russian profits would also decline.

However, in our calibrated model, the gains outweigh the losses in production profits when only the shadow fleet continues to be used after targeting (case (ii)).³⁰ As a result, targeting in case (iii) can either reduce or increase the present value of Russian profits. If we again break the reduction in capacity into the same two steps, Russia gains in the first step (in our calibrated model) and must lose in the second step. If targeting results in a sufficiently small addition of Western services, then in our calibrated simulations it will raise Russia’s present value of profits. On the other hand, if targeting results in a sufficiently small increase in the world price, then it will lower Russia’s present value of profits.

5.2 Simulation results

In the simulations below, we maintain the \$60 cap and simulate unanticipated targeting. In particular, we examine the effect of an unexpected loss in the shadow fleet in the 12th quarter that reduces the size of the shadow fleet to the level Russia had attained 8 quarters earlier: from K_{12} to K_4 .

We examine two cases: In the first, auditing occurs only 20% of the time, so cheating occurs prior to targeting. In the second, the cap is perfectly enforced ($a = 1$), so that no cheating occurs prior to targeting.

In each case, the level of enforcement *after* the 12th-quarter targeting can remain the same or can change: Lax enforcement ($a = .2$) prior to the targeting can be replaced by perfect enforcement afterwards, or perfect enforcement prior to the targeting can be relaxed afterwards.³¹

³⁰For example, with perfect enforcement, if the shadow fleet capacity is reduced from 5.7 mb/d to 5.2 mb/d, this reduction is small enough that only the shadow fleet would be used even after the targeting ($Q = 0$). In this case, targeting causes the present value of profits to *increase* from \$384.3 to \$390.8 billion.

³¹With $a = 0.2$ and $\tau = 10$, capacity that had reached $K_{12} \approx 4.2$ mb/d suddenly drops to $K_4 \approx 2.9$ —a reduction of 31%; with $a = 1$, capacity that had reached $K_{12} \approx 5.7$ mb/d suddenly drops to $K_4 \approx 3.5$ —a reduction of 39%.

Consider first the case where enforcement is initially lax. The solid lines in Figure 7 show the case where enforcement is lax and there is no targeting. The dashed lines show the case in which the lax enforcement level is maintained after the fleet loss. In comparison to the case without a fleet loss, Russian cheating jumps up to offset the fleet loss, so the world price is unaffected. This illustrates case (i). Targeting in case (i) reduces the present value of Russian expected profits. In the bottom right panel of Figure 7, the reduction in the present value of expected profits is 1.87%—from \$380.6 to \$373.5 billion.

In response to the targeting, cheating jumps up but subsequently declines monotonically, ending in the 36th quarter instead of the 28th quarter. Until it ends, the world price does not change. The unexpected destruction of the fleet effectively shifts back the trajectory capacity, eventually converging to the steady state level (top right panel).

We can also simulate what happens if targeting and tightening enforcement occur simultaneously. Suppose we begin with lax enforcement and then simultaneously target and tighten enforcement. We can break this into two steps. If targeting occurred without tighter enforcement, we have just seen that Russian present value of expected profits would fall by a small percentage. If we now tighten enforcement with the fleet capacity held constant, the present value of Russian expected profits increases by a small percentage. As the bottom right panel reflects, the joint effect of targeting and tightening in this case is to lower Russian present value of expected profits imperceptibly (by .05%)—from \$380.6 to \$380.4 billion.

The dotted lines in Figure 7 indicate that switching to perfect enforcement would shut down the cheating channel and would cause the world price of oil to jump up (top left panel). Although aggregate exports jump down, all use of Western services is authorized; Russia no longer cheats because of the increased stringency of enforcement. As a result, the lost capacity is rebuilt in only 8 quarters rather than 28 quarters (top right panel). Consequently, with tightened enforcement, Russian profits are higher merely 3 quarters after the unexpected loss of shadow fleet capacity.

Figure 8 shows the results for the case of perfect enforcement of the cap before the 12th quarter. The solid line reproduces the baseline results. The dotted lines represent the trajectories after the 12th quarter targeting with no change in the enforcement level. This illustrates our case (iii). The top left panel reveals that an unexpected reduction in the shadow fleet causes an upward jump in the world oil price. This loss in fleet capacity induces Russia to resume sales under the

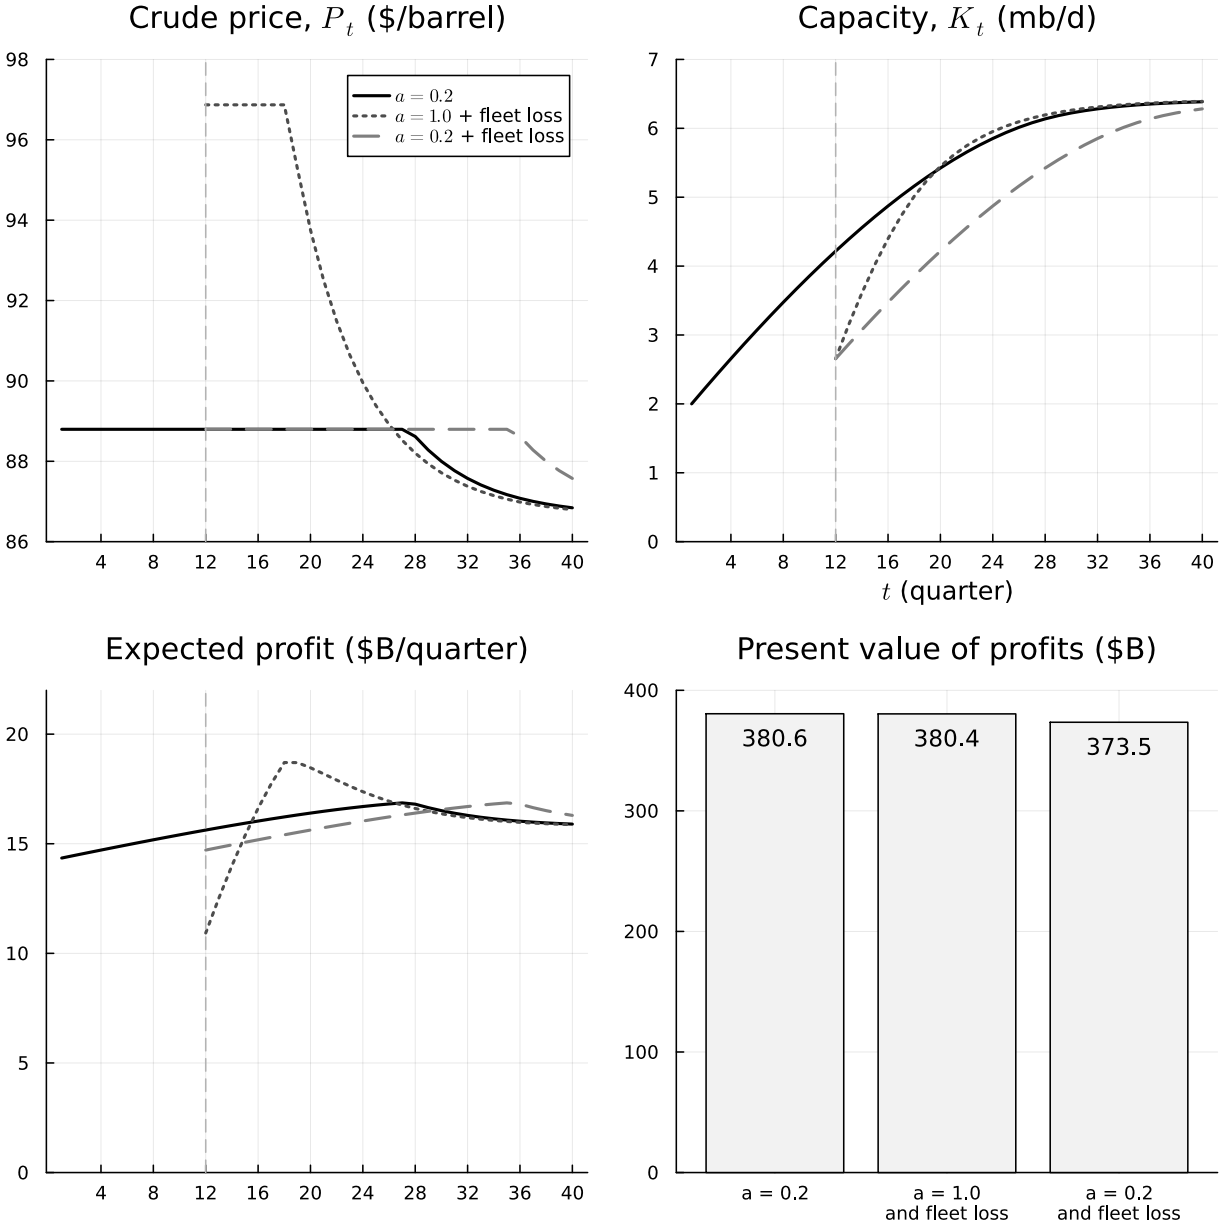


Figure 7: A comparison of the outcomes with and without unexpected changes after the 12th quarter. The solid line represents the low enforcement case ($a = 0.2, \tau = 10$) throughout without shadow fleet loss. The dotted line represents the case with the destruction of the shadow fleet accumulated between the beginning of the 4th and 12th quarters and ramping up to perfect enforcement ($a = 1$) after the beginning of the 12th quarter. The dashed line shows the same shadow fleet destruction but maintaining low enforcement. The line graphs display only the first 40 quarters of the 80-quarter simulation.

cap for the next four quarters. The bottom left panel shows that the unexpected fleet loss causes Russian profits to jump down in the short run when compared to the trajectory without capacity

loss; however, profits after the fleet loss overtake profits had there been no fleet loss after only four quarters. As the bottom right panel reflects, targeting while maintaining perfect enforcement causes the present value of Russian profits to increase by 1.2%.—from \$384.3 to \$389.0 billion.

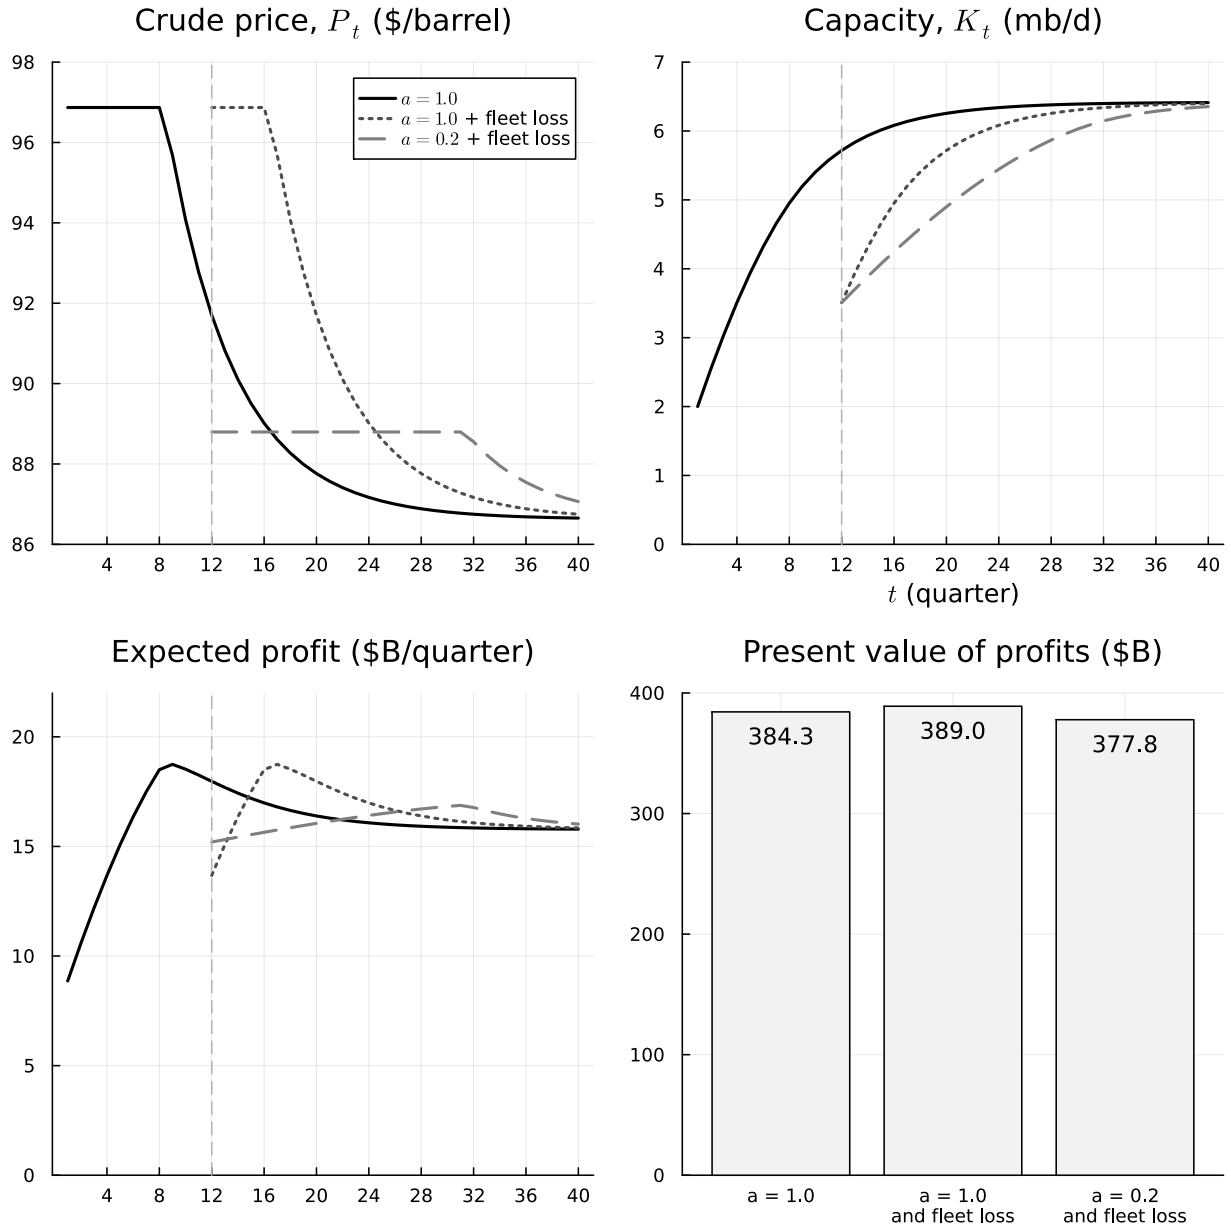


Figure 8: A comparison of the outcomes with and without unexpected changes after the 12th quarter. The solid line represents the perfect enforcement case ($a = 1$) throughout without shadow fleet loss. The dotted line represents the case with the destruction of the shadow fleet accumulated between the beginning of the 4th and 12th quarters while maintaining perfect enforcement. The dashed line shows the same shadow fleet destruction but with low enforcement ($a = 0.2, \tau = 10$) after the beginning of the 12th quarter. The line graphs display only the first 40 quarters of the 80-quarter simulation.

The dashed lines in Figure 8 illustrate the case where enforcement of the cap is relaxed at the same time that the fleet is targeted. As the top left panel reflects, the world price would jump down. Despite the unexpected lower shadow fleet capacity, equilibrium prices are initially lower due to a more than offsetting upward jump in cheating. Since unauthorized sales earn nearly as much as shadow fleet sales, the fleet is rebuilt only slowly. As the bottom right panel reflects, relaxing enforcement when targeting causes the present value of Russian expected profits to fall by 1.7%—from \$384.3 billion to \$377.8 billion.

Suppose auditing is sufficiently frequent that there is no cheating on the \$60 cap. If targeting in the 12th quarter were *combined* with a policy that prevented the rebuilding of the shadow fleet, then the effect on the present value of Russian profits would depend critically on how much capacity was left after targeting.³² In the extreme case where fleet capacity was eliminated entirely, the present value of Russia’s profits over the remaining 69 quarters would drop from \$384.3 billion in the baseline simulation to \$230 billion, a reduction of 40%. In that case, all of Russia’s exports would be sold at the ceiling price. Since there would be no Russian sales at the world price, a tighter cap—although it would elevate the world price—would harm Russia even more, consistent with Hotelling’s lemma.

6 Conclusion

We built and calibrated a model to study the expansion of Russia’s shadow fleet as a response to sanctions against its oil exports. Our model is the first to account for the endogenous, dynamic investment of Russia in its evasion capacity. We find that various policies intended to reduce Russia’s ability to finance the war induce expansion of the shadow fleet with dramatic and unintended consequences.

First and foremost, our analysis suggests that the price cap on Russian oil, although less radical than a complete ban on Western services, is slightly more effective in reducing Russia’s oil profits. In the very short run, a complete ban would have generated sharply lower profits (inclusive of fleet adjustment costs) and a spike in the world oil price; however, the ban would also have precipitated a rapid expansion of the shadow fleet, neutralizing the impact of the sanction. The price cap avoids

³²We are indebted to Catherine Wolfram for suggesting that we analyze the policy combination of targeting and blocking the rebuilding of the shadow fleet.

or mitigates these effects, stabilizing the global oil price and slowing down the expansion of the fleet. Overall, our calibrated simulations show that while both policies significantly reduce the present value of Russia's oil profits, the \$60 cap lowers it slightly more while also benefiting consumers by avoiding a price spike.

Our findings hold further implications regarding the level of the price cap. Under our baseline calibration, the present value of Russian profits is minimized at a cap of \$69.35. Lowering the cap raises the present value of Russian profits modestly and harms consumers by raising the global oil price.

These results are robust to changes in the cost of capacity adjustment, and in Russia's discount rate. The result that both the price cap (at \$60 or lower) and the complete ban reduce Russia's profits in a similar way is also robust to large variations in the elasticity of residual demand for Russia's oil.

The mechanism that makes the price cap more effective also has consequences for the effects of policies currently considered to either strengthen enforcement of the price cap or to shrink the size of the shadow fleet (targeting). We find that both tightening enforcement and targeting prompt Russia to sharply expand its shadow fleet. As a result, these policies could backfire, reducing the loss in the present value of Russia's profits conferred by imposing the cap.

Overall, our results substantiate the arguments in favor of the price cap over the service ban and for maintaining the cap at its current level instead of lowering it. More generally, they also call for increased attention to any policy changes that could induce Russia to expand its fleet more rapidly. Besides, the analysis of our results' sensitivity points to the crucial role of the price elasticity of the residual demand facing Russia, calling for attention to complementary energy policies that would make non-Russian oil production more responsive. Indeed, if the residual demand for Russian oil were more reactive to price changes, existing sanctions would be more effective. For example, if Russian oil export cuts were substantially offset by non-Russian supply increases, spikes in the world price would be dampened, thus lowering the profitability—and expansion—of the shadow fleet.

Our findings are based on a parsimonious model that focuses on essential mechanisms of the policies under consideration. As such, this model might not capture other relevant aspects, a limitation of our analysis. For instance, our paper overlooks Russia's market power in the global oil

market. A more realistic assumption that Russia has some market power on the residual demand that it faces, however, seems unlikely to change our qualitative results significantly. We base this on Appendix A, where we maintain our assumption of unchanging non-Russian production and show that, because of induced shadow fleet expansion, the present value of Russian profits would be even less sensitive to changes in the cap if Russia exerted market power in its most extreme form. In reality, since OPEC also has market power, non-Russian production (Z) is endogenous. Taking account of this more realistic market structure while maintaining our focus on Russia's dynamic expansion of its shadow fleet in response to various Western sanctions will require analysis of a dynamic game, a complex undertaking. We hope to explore this in future work.

Besides policies that have been implemented or actively considered, other reforms of the current sanctions deserve consideration. Most prominently, the price cap mechanism could be extended to include oil products refined in India and elsewhere outside of Russia if they use imported Russian crude oil—such products are currently exempted from the cap. Such reforms are worth careful examination in future modeling.

References

- Babina, T., Hilgenstock, B., Itskhoki, O., Mironov, M., & Ribakova, E. (2023). *Assessing the Impact of International Sanctions on Russian Oil Exports* (SSRN Working Paper No. 4366337). Retrieved from <https://papers.ssrn.com/abstract=4366337> doi: 10.2139/ssrn.4366337
- Becker, R. A. (1985). Capital income taxation and perfect foresight. *Journal of Public Economics*, 26(2), 147–167.
- Becko, J. S. (2024). A Theory of Economic Sanctions as Terms-of-Trade Manipulation. *Journal of International Economics*, 150, 103898.
- Brooks, R., & Harris, B. (2024, June 26). *Why the US should sanction more Russian tankers*. Retrieved from <https://www.brookings.edu/articles/why-the-us-should-sanction-more-russian-tankers/> (Accessed on October 29, 2024)
- Central Bank of Russia. (2024). *Russian Government Bond Zero Coupon Yield Curve, Values (% per annum)*. Retrieved from https://cbr.ru/eng/hd_base/zcyc_params/ (Accessed: 2024-10-16)

- Centre for Research on Energy and Clean Air. (2023). *December 2023 Monthly Analysis on Russian Fossil Fuel Exports and Sanctions*. Retrieved from <https://energyandcleanair.org/december-2023-monthly-analysis-on-russian-fossil-fuel-exports-and-sanctions>
- Gars, J., Spiro, D., & Wachtmeister, H. (2022). The Effect of European Fuel-Tax Cuts on the Oil Income of Russia. *Nature Energy*, 7(10), 989–997. doi: 10.1038/s41560-022-01122-6
- Hilgenstock, B., Ribakova, E., Shapoval, N., Babina, T., Itskhoki, O., & Mironov, M. (2023). Russian Oil Exports Under International Sanctions. *SSRN*. Retrieved from <https://ssrn.com/abstract=4430053>
- Hotelling, H. (1931). The Economics of Exhaustible Resources. *Journal of Political Economy*, 39(2), 137–175.
- International Energy Agency. (2024). *Average Russian Oil Exports by Country and Region, 2021-2023*. Retrieved 2024-10-02, from <https://www.iea.org/data-and-statistics/charts/average-russian-oil-exports-by-country-and-region-2021-2023>
- International Monetary Fund. (2021). *IMF Country Report No. 21/36*. <https://www.elibrary.imf.org/view/journals/002/2021/036/002.2021.issue-036-en.xml>. Retrieved from <https://www.elibrary.imf.org/view/journals/002/2021/036/002.2021.issue-036-en.xml> (Accessed: 2024-10-16)
- Johnson, S., Rachel, L., & Wolfram, C. (2023a). Design and Implementation of the Price Cap on Russian Oil Exports. *Journal of Comparative Economics*. doi: 10.1016/j.jce.2023.06.001
- Johnson, S., Rachel, L., & Wolfram, C. (2023b). *A Theory of Price Caps on Non-Renewable Resources* (NBER Working Paper No. 31347). Retrieved from <https://www.nber.org/papers/w31347> doi: 10.3386/w31347
- Kilian, L. (2022). Understanding the Estimation of Oil Demand and Oil Supply Elasticities. *Energy Economics*, 107, 105844. doi: 10.1016/j.eneco.2022.105844
- Kilian, L., Rapson, D., & Schipper, B. (2024a). *The Impact of the 2022 Oil Embargo and Price Cap on Russian Oil Prices* [Working Paper]. Retrieved 2024-12-22, from <https://sites.google.com/site/lkilian2019/research/working-papers>
- Kilian, L., Rapson, D., & Schipper, B. (2024b). *Response to ‘Understanding the Price Cap on Russian Oil and its Role in Depressing Russian Oil Revenues’* (mimeo). University of California, Davis. Retrieved from <https://rapson.ucdavis.edu/uploads/8/4/7/1/84716372/>

- response_ceepr_critique.pdf (Accessed: 2024-12-22)
- Olson, M. (1962). The Economics of Target Selection for the Combined Bomber Offensive. *Royal United Services Institution Journal*, 107(628), 308–314.
- Olson, M. (1963). *The Economics of the Wartime Shortage: A History of British Food Supplies in the Napoleonic War and in World Wars I and II*. Duke University Press.
- Shatz, H. J., & Reach, C. (2023). *The Cost of the Ukraine War for Russia*. Retrieved from https://www.rand.org/pubs/research_reports/RRA2421-1.html (Accessed: 2024-10-16)
- Spiro, D., Wachtmeister, H., & Gars, J. (2024). *Assessing the Impact of Oil Sanctions on Russia* [SSRN Scholarly Paper]. Retrieved from <https://ssrn.com/abstract=4860148>
- Stokey, N. L., Lucas, R. E., & Prescott, E. C. (1989). *Recursive methods in economic dynamics*. Harvard University Press.
- Turner, D. C., & Sappington, D. E. (2024). On the Design of Price Caps as Sanctions. *International Journal of Industrial Organization*, 97, 103099.
- Wachtmeister, H., Gars, J., & Spiro, D. (2023). *The Price Cap on Russian Oil: A Quantitative Analysis* [SSRN Scholarly Paper]. Retrieved from <https://papers.ssrn.com/abstract=4497485> doi: 10.2139/ssrn.4497485
- Wolfram, C. (2024). *Understanding the Price Cap on Russian Oil and Its Role in Depressing Russian Oil Revenues* (Tech. Rep. No. MIT-CEEPR-RC-2024-05). MIT Center for Energy and Environmental Policy Research. Retrieved from <https://ceepr.mit.edu/wp-content/uploads/2024/08/MIT-CEEPR-RC-2024-05.pdf> (Accessed: 2024-10-22)
- World Bank. (2023). *GDP (current US) - Russian Federation*. Retrieved from <https://data.worldbank.org/indicator/NY.GDP.MKTP.CD?locations=RU> (Accessed: 2023-10-22)
- Yermakov, V. (2024). *Follow the Money: Understanding Russia's Oil and Gas Revenues*. Retrieved from <https://www.oxfordenergy.org/wpcms/wp-content/uploads/2024/03/Follow-the-Money-Russian-Oil.pdf> (Accessed on October 29, 2024)

Online Appendix for

“The Dynamics of Evasion: The Price Cap on Russian Oil Exports
and the Amassing of the Shadow Fleet”

Diego S. Cardoso,^a Stephen W. Salant,^b & Julien Daubanes^c

^aDepartment of Agricultural and Consumer Economics, University of Illinois Urbana-Champaign.

^bDepartment of Economics, University of Michigan, and Resources for the Future (RFF).
Corresponding author: ssalant@umich.edu.

^cDepartment of Technology, Management and Economics, Technical University of Denmark, Center for Energy and Environmental Policy Research, Massachusetts Institute of Technology, and CESifo.

A Unchanging monopoly behavior in response to changes in the price cap

A paradoxical result in our dynamic model is that Russia, a profit-maximizer, can benefit if the government tightens the price cap even though Russia could have secured the same benefits without any change in the cap. This result is not an artifact of our price-taking assumption. Turner and Sappington (2024) reach the same conclusion in their static Cournot duopoly. As mentioned in the introduction, such a paradox cannot arise if Russia unilaterally controls the world price. In this appendix, we examine this case by assuming that Russia acts like a monopolist facing constant output (Z) from the rest of the world’s oil producers.

We have two objectives: (1) to simulate the boundaries of the present value of Russian profits as a function of the cap for the extreme cases of monopoly and price-taking; and (2) to show the invariance of the present value of Russian profits in the monopoly case as the cap is lowered.

We modify the problem in Section 2.3 by replacing $\{p_t\}_{t=1}^T$ with $p(Q_t + X_t + Z)$. As a consequence, the first-order conditions (1) and (2) are modified as follows

$$Q_t \geq 0, \quad \hat{p} + X_t p'(Q_t + X_t + Z) - C'(Q_t + X_t) \leq 0, \quad c.s. \quad (1)$$

$$X_t \geq 0, \quad p(Q_t + X_t + Z) - d + X_t p'(Q_t + X_t + Z) - C'(Q_t + X_t) - \alpha_t \leq 0, \quad c.s. \quad (2)$$

The remaining equations, (3) and (4), are unchanged. The second term in condition (1) reflects the fact that a barrel sold at the cap depresses the world price just as much as a barrel sold at the market price.

We simulate the solutions of both the price taker and monopolist for cap levels between \$30 and \$70 dollars per barrel.¹ All simulations use our baseline calibrations of the cost functions C and F . This exercise assumes that the monopolist behavior starts at $t = 1$.

The dashed line in Figure A1 shows the present value of profits Russia earns as a price setter responding to different price caps. As expected, the present value of Russian profits is higher under monopoly than under price taking. Note that the present value of profits no longer increases as the

¹For cap values above \$71, the shadow fleet is never used in the competitive solution. For caps below \$34.1/barrel, the Western services are never used—just as if there had been a service ban. This lower bound corresponds to $C'(\bar{K}_1) \approx 34.1$, i.e., the marginal cost associated with allocating all production to the initial shadow fleet capacity.




Figure A1: A comparison of profits under price-taking and monopolist behavior under various cap levels. Values are reported in billions of dollars in present value using the baseline discount rate (15% per year) until the sanction termination (80 quarters).

cap decreases.

In contrast, the solid line in Figure A1 shows that, under price-taking, the present value of profits can decrease as the cap is tightened. Raising the cap above \$34 decreases the present value of profits, with the global minimum attained at a cap of \$69.35 per barrel. Granted, the effect on the present value of Russian profits of changing the level of the cap is very modest—but this itself is an important finding since it suggests that further heated and divisive debates about the level of the cap are unnecessary.

Figure A2 displays the trajectories of prices, profits, capacity, and investment if Russia acted like a monopolist; the lines representing the price-taker solutions are the same as in Figure 2. The two monopoly trajectories in each of the four panels for the complete service ban and the \$60 cap are *identical* and coincide. The only exception is the price (upper left panel) in the very first quarter. Being unable to enlarge 1st quarter shadow fleet capacity, a monopolist would sell a small portion of exports using Western services ($Q_1 \approx 0.7$ mb/d). However, after the 1st quarter, the monopolist never uses Western services again, and the price trajectory thereafter coincides with the monopolist's response to a service ban. For both levels of the cap, the monopolist expands the shadow fleet far less than a price-taker would. As a result, the world price and the monopolist's




Figure A2: A comparison of the trajectories for equilibrium price taking and monopoly behavior when the sanction is a high price cap (\$60) or a low price cap (equivalent to a service ban). For the monopolist, the two trajectories in each panel are indistinguishable since they *coincide*. Each panel displays quarters 1-40 of the 80-quarter simulation.

profit are uniformly higher.

B Sensitivity analysis of the baseline model

In this appendix, we perform sensitivity analyses on key model parameters to examine the impact of different values on the model’s qualitative findings. In particular, we examine how model outcomes vary under different price elasticities of demand and non-Russian supply, discount rates, and marginal investment costs, and whether those changes affect our findings.

Price elasticity of demand

The price elasticity of world demand for oil (ϵ_D) plays an important role in determining price levels in response to shifts in Russian exports. To assess how this parameter affects model outcomes, we consider the absolute value of price elasticities of demand (ϵ_D) varying between 0.05 and 0.75—six times larger than the baseline elasticity of -0.125 . The curves in Figure B1 compare the present value of Russian profits under a service ban and a \$60 cap for a wide range of elasticities.

Based on Figure B1 we reach two conclusions. First, we note that the difference in profits between the cap and ban policies is affected by the choice of the price elasticity of demand—and may even flip sign. Lowering the absolute value of ϵ_D towards zero rapidly amplifies the difference, which is a consequence of the profits under the ban increasing faster than under the cap. Moving in the opposite direction, increasing the magnitude of the price elasticity can change the sign of the difference in profits: Roughly doubling the baseline elasticity leads to a switch in the ordering, with the cap policy delivering slightly larger present value of profits than then service ban.²

The second key message from Figure B1 is that substantial increases in $|\epsilon_D|$ only slightly increase the differences in profits. We observe that as the magnitude of the elasticity grows even beyond typical empirical estimates, the curves flatten out, so that the gap between profits in each case remains stable. Therefore, even though a more elastic demand would flip the ordering of profits under the service ban or the cap, both policies continue to result in very similar profits in present value.

Price elasticity of non-Russian supply

In the baseline version of our model, we have assumed that non-Russian sources of oil are constant (Z). Had we assumed, instead, that oil supply from the rest of the world was sensitive to the oil price ($Z(p)$), changes in the price elasticity of this supply would be analogous to changes in the price elasticity of demand for Russian oil $|\epsilon_D|$. It follows that the previous analysis of how the price elasticity of demand affects our results captures the sensitivity of our results to changes in non-Russian supply responses. In particular, an upward-sloped non-Russian supply might flip the ordering of profits under the cap or the service ban, but would not change the fact that both sanctions deliver quantitatively similar profits in present value.

Discount rate

Next, our analysis considers discount rates from 0 to 30% per year—twice as high as the baseline rate of 15%. Figure B2 shows how the difference between the present value of profits between the cap and the ban policies varies with different combinations of the discount rate and elasticity.

²For reference, empirical estimates of this elasticity typically range between 0 and -0.4 , depending on the estimation approach, time horizon, and data used. See Kilian (2022) for a detailed discussion of the estimation of price elasticities in the oil sector.




Figure B1: Cumulative profits under different policy scenarios for a range of price elasticities of demand. Values are reported in billions of dollars in present value using the baseline discount rate (15% per year) until the sanction termination (80 quarters).

Negative values indicate that the PV of profits is lower under the baseline cap (\$60) relative to a ban policy.³

The almost vertical level curves in Figure B2 demonstrate that the ordering in the outcomes between the cap and ban policies is largely insensitive to the choice of the discount rate. This result is primarily due to the calibration of the linear marginal investment cost function F' . Since we fit the slope of this function to match observed data, an increase in the discount rate leads to a lower marginal investment cost so as to match the observed level of investment.

To further illustrate how the equilibrium paths under competition are fairly insensitive to discount rate choices, Figure B3 shows a panel of trajectories for the cases of halving and doubling the baseline rate. These panels show that trajectories under the cap policy vary little, and those under the service ban are essentially the same. Therefore, our qualitative baseline results remain robust within a reasonable range of discount rates.

³The simulations reported in Figure B2 use different discount rates for Russia. However, the displayed present value of profits is calculated using a common rate to ensure fair numerical comparisons across different scenarios.




Figure B2: Difference in profits between the cap and the ban policies for various combinations of the discount rate and the price elasticity of demand. Values are reported in billions of dollars in present value discounted at the baseline discount rate (15% per year) until the sanction termination (80 quarters).




Figure B3: A comparison of the trajectories under a price cap sanction (black) and a service ban (gray). Solid lines represent trajectories simulated using a 7.5% annual discount rate, whereas dotted lines represent these trajectories using a 30% discount rate. All scenarios are based on the baseline price elasticity of demand ($\epsilon_D = -0.125$). These panels display only the first 40 quarters of the 80-quarter simulation.

Marginal investment cost

Next, we consider the impact of changing the parameter ϕ that determines the marginal cost of expanding the shadow fleet. In our simulations, we calibrate ϕ so that the solution of the baseline model (with a \$60 cap) replicates the observed expansion of exports to non-Western countries. In this analysis, we vary ϕ to assess its effect on the differences in profits between the cap and service ban policies.

In Figure B4, we examine how different values of ϕ affect simulation outcomes. The medium ϕ scenario reproduces the baseline simulation, with a calibrated $\phi = 4.102$. The other two scenarios, low and high ϕ adopt respectively half and double the calibrated ϕ . These panels show that a higher marginal investment cost leads to a slower expansion of the shadow fleet in either sanction. Therefore, a higher ϕ corresponds to a longer price plateau and a later peak in profits. However, as the bottom right panel shows, the present value of profits is higher under the service ban relative to a \$60 cap regardless of the value of ϕ .

Figure B5 examines whether the relative ranking of profits under a service ban or cap flips with higher values of ϕ . As this figure indicates, the curves representing profits under each sanction do not cross even for values of ϕ that are 10 times larger than the calibrated one. Moreover, we observe that the gap between profits increases: A higher ϕ leads to higher world prices under the service ban, which boosts short-run profits relative to a longer plateau seen under the cap.

Discussion of alternative functional forms for the investment cost. Besides changing the slope of the marginal investment cost function F' , one could also consider two additional changes. First, we could introduce a positive intercept so that $F'(0) > 0$. The practical result of a positive intercept would be that the size of the shadow would converge to a lower value, as investment levels would stop short of the point where $C'(K^*) = P(Z + K^*) - d$. Second, we could make F a function of K_t to introduce dynamic convexity in investment costs. Again, this modification would result in the shadow fleet converging to a value lower than K^* . Although such modifications are technically possible, we opt for a parsimonious definition of F for two reasons: (i) The shadow fleet expansion is still in progress as we write this paper, so we do not have reliable data to calibrate parameters governing the convergence point, and (ii) our adopted F allows us transparently to link the final size of the shadow fleet with calibrated marginal production costs and producer prices, as

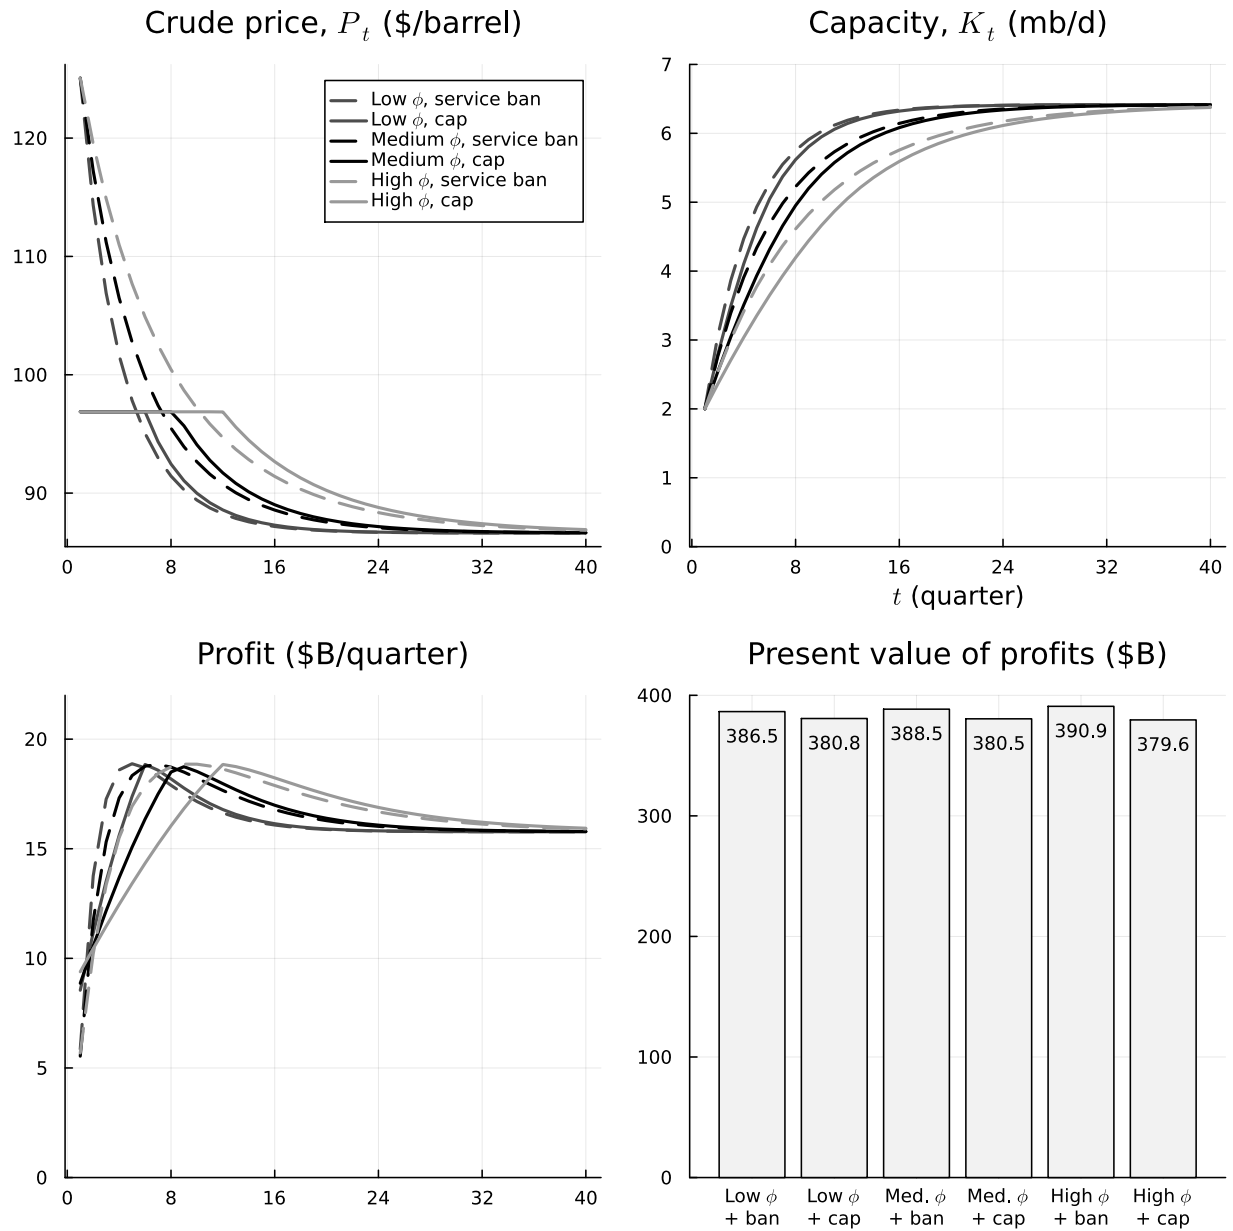


Figure B4: A comparison of prices, capacity, and profits under the service ban (dashed lines) vs. a \$60 price cap sanction (solid) with different marginal expansion costs (ϕ). The line graphs display only the first 40 quarters of the 80-quarter simulation.

represented in the diagram in Figure 1.

References

Kilian, L. (2022). Understanding the Estimation of Oil Demand and Oil Supply Elasticities. *Energy*

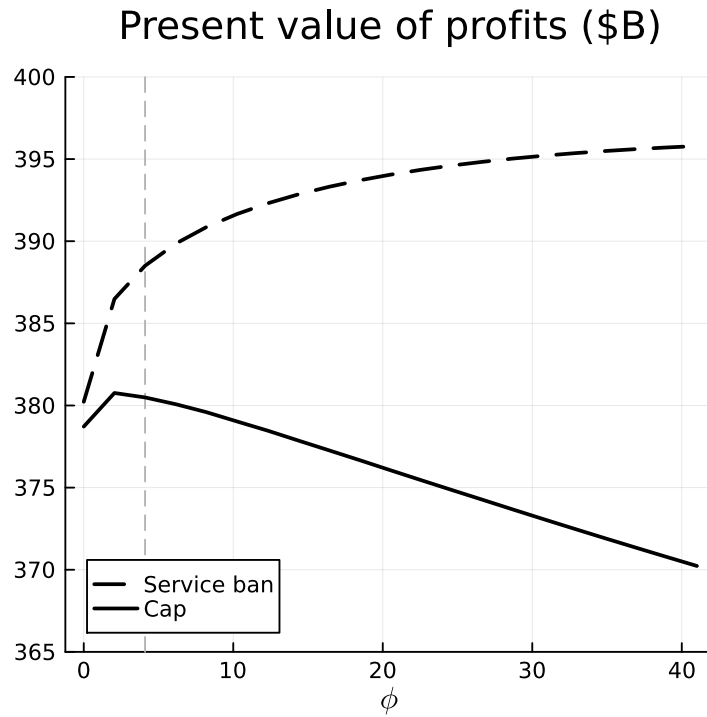


Figure B5: Present value of profits under the service ban or \$60 price cap sanction for different levels of the marginal investment cost (ϕ parameter).

Economics, 107, 105844. doi: 10.1016/j.eneco.2022.105844

Turner, D. C., & Sappington, D. E. (2024). On the Design of Price Caps as Sanctions. *International Journal of Industrial Organization*, 97, 103099.