ECONOMISTS FOR UKRAINE

WORKING PAPER SERIES

WORKING PAPER No. 5

SEPTEMBER 2025

THE DYNAMICS OF EVASION: THE PRICE CAP ON RUSSIAN OIL EXPORTS AND THE AMASSING OF THE SHADOW FLEET

Diego S. Cardoso, Stephen W. Salant, and Julien Daubanes

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

Economists for Ukraine (Econ4UA)

Website: https://econ4ua.org/ Email: info@econ4ua.org

ABSTRACT

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

To reduce Russia's funding for its Ukraine invasion, Western governments imposed, after a delay, a price ceiling on Russian seaborne oil exports utilizing Western services. To evade that ceiling, Russia developed a "shadow fleet'' using no such services. We simulate a calibrated model driven by this fleet's expansion to assess various sanctions. Mainly, sanctions—from a price ceiling to its extreme service ban version—significantly reduce the present value of Russia's profits. However, tighter price caps will not necessarily harm Russia if they raise the world price. Shortening the delay between sanctions announcement and implementation can harm Russia more.

JEL CLASSIFICATION: D04, L51, Q41
KEYWORDS: Economic warfare, Sanctions evasion, Hotelling's lemma

Diego S. Cardoso
University of Illinois Urbana-Champaign
Department of Agricultural and Consumer Economics
Champaign, IL
dcardoso@illinois.edu

Stephen W. Salant
University of Michigan
Department of Economics
& Resources for the Future
Ann Arbor, MI
ssalant@umich.edu

Julien Daubanes
Technical University of Denmark
Department of Technology, Management and Economics
& CESifo
Kongens Lyngby, Denmark
jxada@dtu.dk

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. At the same time, 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 went to pay for the Ukraine war.[1] Given this situation, Shatz and Reach (2023) concluded: "Blocking those [revenues] could be the most powerful tool in the West's economic toolkit to hamper Russia's war effort ..."

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 June 2022, the European Union (EU) finalized its agreement to ban the use of Western services on all Russian seaborne exports. The policy was to go into effect six months later.

As global oil demand is inelastic and Russia was producing about 11% of the world's oil, however, 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 G-7, 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.

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 has assembled a fleet of aging tankers (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

[1] Our 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).

[2] Western 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."

enforcement of the cap has also led Russia to cheat occasionally by selling oil at prices above the cap while using Western services. Violators are subject to punishment. Recently, Western countries have implemented new policies to deter such cheating and to reduce the size of the shadow fleet—for example, a significant expansion of penalties against the shadow fleet, as advocated by Brooks and Harris (2024), was enacted by the US Administration in January 2025.[3] 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 adjustments and policy-makers currently considering a tighter price cap.[4] There have even been suggestions to apply similar sanctions to Iran and other countries.[5] 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, as our subsequent literature review reflects. 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 Russia's shadow fleet. We then calibrate model parameters to reproduce observed facts in data and use it to simulate the outcomes of various policies intended to reduce Russian oil profits.[6] We say "intended"

[3]See, e.g., `https://home.treasury.gov/news/press-releases/jy2777` for more details.

[4]See, e.g., Simon Johnson and Oleg Ustenko's suggestion that the 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.

On G-7 recent 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` and *The Financial Times*, May 27, 2025, available at `https://www.ft.com/content/ef479183-32ca-486e-9dea-ff17c4280922`.

[5]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` has been 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.som.yale.edu/insights/to-prevent-wider-war-in-the-middle-east-choke-off-irans-oil-sales`.

[6]We focus on profits from oil exports both because they measure exports' value added to the overall Russian economy and because the Russian government's take from its oil industry has become increasingly oriented towards the taxation of profits, a trend we expect to continue over the horizon in our simulations. Over the past few years, the tax system has moved away from the export tax—which was due to be completely abandoned by 2024—and towards a profit tax system (Yermakov, 2024). Moreover, since the Russian oil industry is mostly owned by the State, residual profits accrue to the government through dividends. Using Rystad Energy data, Spiro et al. (2024) estimate the Russian government's take to lie between 70% and 80% of the industry's total profits.

because we have discovered that under our baseline assumptions, some policies designed to shrink Russian profits do lower profits (producer surplus net of the cost of expanding the shadow fleet) in the very short run but, in fact, raise future profits so much that the present value of the entire stream of profits increases. We also examine how these results are affected if we relax our baseline assumptions.

In our dynamic model, Russia maximizes its discounted profits from oil exports. 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 discount using the shadow fleet—this discount is in part due to Western countries' embargo, which underlies each of the policies we consider in this paper.[7] The world price path adjusts to clear the market. We assume that Russia takes the world oil price path as given, consistent with Cahill's (2023) account of Russia's supply behavior,[8] although we examine the consequences for our results of Russia's internalization of its market influence in Subsection 6.4. To isolate the effects of shadow fleet expansion, our simulations take no account of the depletion of Russian oil reserves or the possibility of aboveground storage. 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. In our baseline simulations, we also assume that non-Russian sources of supply are constant but explore the effects of relaxing this assumption in Subsection 6.1.

We assume the cap is binding and sufficiently tight that Russia always utilizes the shadow fleet. Our baseline analysis assumes sanctions were implemented immediately after the Ukraine invasion, but we explore the effects of anticipation of delayed sanctions in Subsection 6.5. Initially, Russia has limited shadow-fleet capacity because, prior to the sanctions, Russia relied on Western services. Given the initial limitation on fleet sales, the marginal cost of additional production is below the \$60 cap, so if the cap is implemented immediately, 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

[7]The Western embargo, implemented around the same time as the \$60 price cap, is one factor underlying the discount on Russian shadow-fleet exports. One implication of this discount is that non-coalition countries—such as China and India, which opposed the sanctions—imported oil at a price lower than the global oil price.

[8]According to Henderson and Fattouh (2016), "Russia has historically adopted a consistent line that it is a price taker when it comes to the oil market." They partially attribute Russia's apparent price-taking behavior to the hybrid structure of its oil industry in which oil supply decisions are largely decentralized.

imposition of the cap.[9] There are two phases to the world price path. In the first phase, the world price is predicted to remain constant since the growth of fleet sales is exactly offset by the decline in ceiling sales. In the second phase, the world price falls because, once cap sales cease altogether, fleet sales continue to expand.

We use our model to compare the effects of reducing the ceiling price or, in the extreme case, setting it so low that it is equivalent to the service ban proposed by the EU. Russian profits (defined as producer surplus minus adjustment costs) under the ban would at first be considerably less than under the \$60 cap because 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 through the 8^{th} 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 ($\approx 2\%$) and can be reversed with a plausible relaxation of our baseline assumptions. 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.

Russia maximizes the present value of its profits when facing a binding price ceiling but, under our baseline assumptions, would earn a higher present value if it instead faced a tighter cap or, in the extreme, a service ban. This result may seem paradoxical, but it can arise in equilibrium problems where the maximizer takes as given a variable he does not control—the world price in the case of competitive equilibrium and the sales of a rival in the case of Cournot equilibrium. The paradox rests on reasoning which, although seemingly intuitive, is fallacious (except in single-agent problems) that an agent can duplicate the outcome of an equilibrium by duplicating his own actions. But when a new government policy is announced, the behavior of the maximizing agent is not the only thing that changes; the variable that the maximizing agent has taken as given also changes. That is why, for example, the home exporter in Brander and Spencer's (1985) Cournot duopoly game benefits from a subsidy even if the subsidy payments are removed lump sum. The export subsidy not only induces the home firm to expand its exports; it also causes the exports of the foreign firm, which the home firm had taken as given, to contract. Similarly, a lower price cap

[9] We return to this prediction of a price jump in Subsection 6.5, where we extend the model to allow anticipatory behavior.

not only causes Russia's exports to contract; but it also causes the world price, which Russia had taken as given, to increase.

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

Our simulations indicate that the rate of expansion of the shadow fleet is relatively sensitive to the cap level. In the interval between \$34 and \$69, lowering the cap slightly *raises* the present value of Russian profits 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 1^{st} quarter, 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 1^{st} 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. A cap of \$69.35 minimizes the present value of Russian profits.

At the end of Section 4, we provide intuition for these dynamic results by showing graphically and analytically how tightening the price cap creates opposing effects on the producer surplus in any given quarter given Russia's current evasion capacity. Under our baseline assumptions, these opposing effects result in a small increase in producer surplus underlying the increase in the present value of Russian profits (producer surplus net of capacity adjustment costs). When enforcement is imperfect, this same static intuition suggests why increased enforcement may raise the present value of Russian profits. In Section 6, we relax these baseline assumptions and identify the magnitude of various parameters sufficient for tightening the cap to *lower* the present value of Russian profits. The assumed 1^{st}-quarter fleet size is calibrated based on pre-invasion trade flows. Although conservative, it is nonetheless sufficiently large that reducing the cap (or imposing a service ban) always increases the present value of Russian profits.

Next, we investigate the effects of policies affecting the shadow fleet. We begin by determining the value to Russia of having a shadow fleet and the value to Russia of having the flexibility to expand it. To determine this, we simulate the consequences of Russia having a fleet of a given initial size (either 0 or 2 mb/d) but being unable to expand it. This also tells us, from the West's perspective, the benefit the West would have reaped if it had prevented any fleet expansion. We then consider the consequences of policies that steepen the linear marginal cost of expanding the fleet. Finally, we simulate an unanticipated reduction of the shadow fleet in the 12^{th} quarter, which lowers the fleet size to where it had been in the 4^{th} 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, in our simulations using baseline calibration, 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, Rachel, and Wolfram (2025). 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 points to the usefulness of the price cap in reducing Russian oil profits and in stabilizing oil prices. It also highlights that the existence of a shadow fleet, which allows Russia to evade the cap, can undermine the effectiveness of the instruments. Yet, they leave the expansion of this fleet unexplored. Although our models are very different,[10] the main novelty of our framework of analysis is the endogenous and dynamic development of the fleet capacity.

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 (2023) examine the effect of price caps lower than \$60 in a static setting. Like Johnson et al. (2025), 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 and would

[10] A major difference is that we do not treat oil as an exhaustible resource. For simplicity, as well as for calibration purposes, we assume a stationary supply curve for oil. The adoption of Johnson et al.'s (2025) oil supply model in our parsimonious setup would substantially complicate the analysis at the expense of transparency on dynamic mechanisms.

have generated significantly higher global oil prices. Neither paper takes into account how lower caps hasten the expansion of the shadow fleet and how this affects Russian profits.

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. This increase in aggregate Russian exports causes a decrease in non-Russian exports due to strategic substitutes. Since the induced decrease is less than offsetting, the world price falls. As for Russia's profits, they increase compared to *laissez-faire* exactly like those of the home exporter in Brander-Spencer's Cournot game (Brander and Spencer, 1985).[11] Our results are completely different. In our model, the imposition of the import embargo and the $60 cap reduces Russia's sales using Western services, raises the world price, and reduces the present value of Russian profits by approximately 25% compared to *laissez-faire*.

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 using the fleet to evade the price cap requires Russia to incur various costs, including increased cost of shipping longer distances (e.g., to India, as a result of the EU embargo), heightened premia on insurance against oil spills, and strengthened monopsony power of buyers of Russian oil. Johnson and Wolfram (2024) also suggest that discounts on Russian shadow-fleet exports are due to the premia on insurance to cover the risk of direct sanctions on owners of tankers caught using Western services while carrying oil priced above the cap. In our model, we consolidate these costs into an exogenous fixed cost per barrel shipped. We explore the consequences of endogenizing this cost in Subsection 6.3.

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, the effects of tightening the cap, increasing enforcement, and sanctioning the shadow fleet in various ways.

[11] In both duopoly models, these results are reversed if the policy is sufficiently extreme. Turner and Sappington (2024) regard the caps below $56.35 as outside the relevant region, but in that region revenue, profits, and sales using Western services are strictly increasing in the cap level. A very low cap can reduce Russia's profits by up to 19%.

To put this literature on sanctioning Russia's oil exports in perspective, our analysis illustrates Mancur Olson's (1962, 1963) explanation for why targeted economic sanctions often have a limited impact. Analysts often identify inputs that—before the commencement of hostilities—an adversary was relying upon as if essential, only to discover once the war starts and these inputs are targeted that they can be readily replaced by substitutes.[12]

We proceed as follows. In Section 2, we describe our model. In Section 3, we describe its baseline calibration. In Section 4, we use our calibrated model to compare the effectiveness of price caps set at various levels including at the level equivalent to a service ban. We also investigate the effects of increasing the level of enforcement of the cap. In Section 5, we consider various policies directly targeting Russia's shadow fleet. In Section 6, we discuss the consequences of relaxing our baseline assumptions. Section 7 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.[13]

Russia sells the rest of its oil at a price higher than the cap using its 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

[12] During World War 2, the young economists in the Enemy Objectives Unit (EOU) of OSS (including Carl Kaysen (22 years old), Harold Barnett (25), Walt W. Rostow (26), William A. Salant (26), Charles Kindleberger (32), and Emile Despres (33)) recognized that Allied resources were limited and developed criteria for determining the optimal German targets to bomb. The ability of the Germans to substitute was a chief concern. W. A. Salant, who drafted the first statement of EOU's theoretical framework, wrote: "It is easy to imagine a material which is used entirely for direct military purposes, whose supply cannot be expanded, but for which a satisfactory substitute is readily available. Because of the possibility of easy substitution this industry would have to be discarded as a potential target ..." (Salant, 1942). As their modern-day counterparts are discovering when devising sanctions against Russia, applying these theoretical criteria sometimes proves difficult in practice. In World War 2, ball bearings were deemed essential to German munitions production. Consequently, the Allies bombed several ball-bearing factories. In response, the Germans substituted by rebuilding the bombed facilities, expanding production at unbombed ball-bearing factories, importing ball bearings from Switzerland and Sweden, de-cumulating ball bearing inventory, and using slide bearings, etc. Munitions production was barely affected.

[13] We 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 complemented 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.

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. When enforcement of the cap is imperfect, an additional cost of using the fleet is the risk of direct sanctions on owners of fleet tankers caught cheating—such as those the US Administration enacted in January 2025.

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. We endogenize this discount in Subsection 6.3.[14]

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.[15] We assume neither the demand function nor non-Russian supply fluctuates over time. We assume that oil is inexhaustible and that there is no aboveground storage capacity in Russia or elsewhere.[16] 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 these costs, Russia must also incur additional costs such as the cost of port expansion and improvement. We model the fleet as capital and its expansion as a costly investment.

Finally, in order to exclude any anticipatory effect, we assume that sanctions were implemented immediately after the Ukraine invasion. This eliminates expansion of the shadow fleet in the months during that period. We extend our analysis to delayed sanctions and to anticipatory behavior in Subsection 6.5.

2.2 Notation and Curvature of Functions

We adopt the following notation:

- K_t : the capacity of the Russian shadow fleet in period t.

[14]Although the monopsony power of importers of Russian oil delivered using the shadow fleet is often mentioned, we do not model it explicitly in our baseline analysis; Kilian et al. (2024a) suggest that it accounts for a transitory part of the total cost of evading sanctions through the shadow fleet. Nor do we explicitly model the risk premium resulting from sanctions on fleet operators when we consider enforcement is imperfect in Subsection 4.3. In Subsection 6.3, however, we discuss how our results are affected by extending the baseline model to introduce an endogenous discount.

[15]In Appendix B.4 we assume instead that Russia takes account of its own effect on the world price.

[16]According to Johnson et al. (2025), existing possibilities of oil storage in Russia are either quantitatively insignificant or economically irrelevant.

- $\bar{k}$: 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_0 : Non-Russian sales in each period.

- $C(R_t)$: Russian total cost of producing oil 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 $I_t > 0$.

- $P(Z_0 + R_t)$: world inverse demand for oil evaluated in period t ($P'(Z_0 + R_t) < 0$).

- d : price discount per barrel on shadow fleet sales.

- $\hat{p}$: price cap.

- β : constant discount factor.

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 $t = 1, \ldots, T$. 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 $t = 2, \ldots, T$, 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:

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

The following conditions must hold with complementary slackness (abbreviated as c.s.):[17]

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

$$X_t \geq 0, \qquad (p_t - d) - C'(Q_t + X_t) - \alpha_t \leq 0, \ c.s.; \tag{2}$$

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

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

To deduce the conditions holding in the competitive equilibrium, we assume the market clears in each period by replacing p_t by $P(Z_0 + Q_t + X_t)$.[18]

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

[17] That 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.

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

$$\text{where } H(K_t) = \begin{cases} K_t[P(Z_0 + R_c) - d] & \text{for } K_t < R_c \\ R_c[P(Z_0 + R_c) - d] + \int_{R_c}^{K_t} [P(Z_0 + u) - d] du & \text{for } K_t \geq R_c \end{cases} \tag{5}$$

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 equilibrium distorted by government intervention.

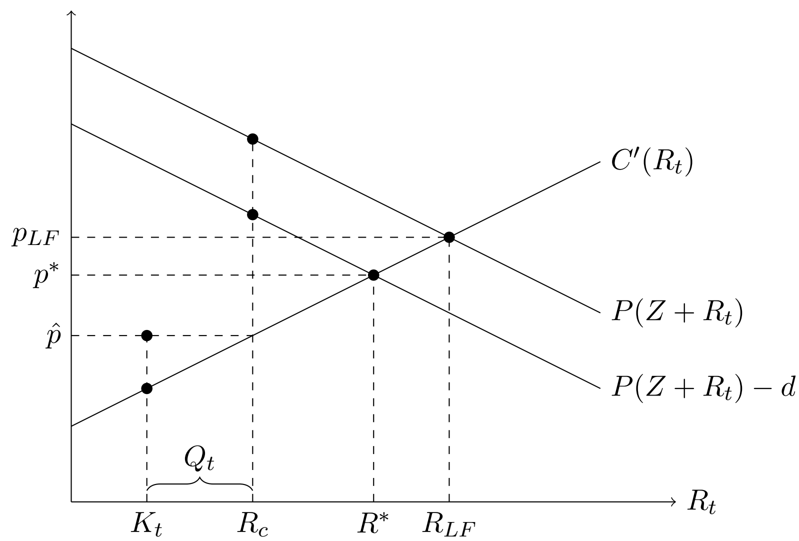

Figure 1: Equilibrium depends on whether cap is non-binding ($\hat{p} > p_{LF}$), slightly binding ($\hat{p} \in [p^*, p_{LF}]$), or tightly binding ($\hat{p} \in [C'(0), p^*)$).

of Russian aggregate sales, $P(Z_0 + R_t)$. The lower downward-sloping curve is this inverse demand curve shifted down by the exogenous constant discount, $P(Z_0 + R_t) - d$. 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, $\hat{p} \in (p^*, p_{LF})$.[19]

Since the price in the two cases strictly exceeds what Russia would receive using the 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 intersection of the net inverse demand function and the marginal cost curve. There are two possibilities: (i) $K_t \leq R^*$ and (ii) $K_t > R^*$. In the first case, $X_t = K_t$ and $\alpha_t \geq 0$. In the second case, $K_t > X_t$ and $\alpha_t = 0$. In the second case, Russia would not expand its fleet and, in fact, would not use all of it. In our calibration,

[19]The "slightly binding" case disappears if $d = 0$. In that more conventional situation covered in the textbooks, the cap either binds or does not bind.

$K_1 = \bar{k} < R^*$. In that case, the second case never arises.

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 $P(Z_0 + R_c)$.

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 $P(Z_0 + R_c)$. 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 $P(Z_0 + R) - d = C'(R)$ and $p^* = P(Z_0 + R^*)$. 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

We specify parameter values and functional forms to make the model numerically tractable. This section outlines the specifications and rationale for our choices. Key calibrated parameters are presented in Table 1. We consider multiple extensions to relax assumptions of the baseline model in Section 6 and perform sensitivity analyses of parameters in Appendix A.

In our model, time periods are quarters, and all volumes are numerically aggregated to quarters. However, in this paper, we report volumes in million barrels per day (mb/d) to stay consistent with the unit typically used in oil markets. Our baseline simulations adopt 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).

Parameter	Variable	Value	Source
Russian time horizon	T	80 quarters	Modeling assumption
Discount factor	β	0.960	Modeling assumption
Non-Russian supply	Z	92.0 mb/d	IEA, Gars et al. (2022)
Initial capacity	$\bar{k}$	2.0 mb/d	IEA, modeling assumptions
Pre-invasion supply to 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 production cost intercept	c_0	\$17	Rystad, Gars et al. (2022)
Marginal production cost slope	γ	0.095	Calibrated to match IEA data
Marginal investment cost 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—as the "non-coalition" regions and set the initial evasion capacity equal to the daily average volume of Russian exports to these regions in 2021: $\bar{k} = 2.0$.[20] Consequently, initial export volume to all other regions, or the "coalition," is $Q_0 = 5.4$ mb/d. Since the baseline model assumes sanctions come into effect at $t = 1$, this implies that simulations reflect the implementation of sanctions immediately after the invasion.

As in Gars, Spiro, and Wachtmeister (2022), we use IEA data to determine the non-Russian global supply of oil, with $Z_0 = 92.0$ mb/d assumed constant throughout. Moreover, we adopt a pre-invasion reference crude price of $p_0 = \$80$ per barrel, approximately reflecting average Brent prices in the 4^{th} quarter of 2021 (International Energy Agency, 2022). 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 any war-related sanctions. We define BAU as the repetition of pre-invasion prices and quantities—which we denote by $t = 0$—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

[20]The daily average volumes are, respectively, 1.6. 0.1, 0.2, and 0.1 mb/d. We assume these regions as the "non-coalition" based on the observation that these regions showed early evidence of receiving Russian exports via shadow fleet shipments and increased their oil imports from Russia post-invasion. In contrast, the EU, United Kingdom, US, Japan, and South Korea jointly accounted for 4.4 mb/d in 2021, while all other regions combined for 1.0 mb/d.

stabilized one year after the start of sanctions. This value is in line with Kilian et al. (2024a), who find that increased shipping distances led 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 et al., 2022; Wachtmeister et al., 2023; Johnson et al., 2023, 2025).

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

$$P(Z_0 + Q_t + K_t) = p_0 \left(\frac{Z_0 + Q_t + K_t}{Z_0 + Q_0 + \bar{k}} \right)^{1/\epsilon_D} \tag{6}$$

where ϵ_D is the price elasticity of demand. Our baseline case in the simulations follows Gars et al. (2022) and adopts a price elasticity $\epsilon_D = -0.125$; this value is supported by recent empirical estimates (see Kilian (2022) for a review). 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_0 of it, the magnitude of the elasticity of the residual demand rises with price and residual demand is zero at prices above approximately \$150 ($P(Z_0) \approx \$150$).

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 et al. (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 1^{st} 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 at \$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. This method searches for the parameter value in which the solution of the model gives a shadow fleet capacity at the end of the 8^{th} quarter (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 $\phi \approx 4.1$.[21] When performing sensitivity analyses of other parameters and extending the model to relax baseline assumptions, we recalibrate ϕ following the same procedure to ensure comparability across results.

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

This section consists of four subsections. The first one shows our main result: The present-value of Russian profits is largely reduced by the sanctions, regardless of whether the \$60 cap is imposed or whether 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, despite the fact that the latter is a more radical policy.

This result is substantiated by the second and third subsections, respectively examining the policy of lowering the cap below its current level, as many supporters of Ukraine advocate, and relaxing our assumption of perfect 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 provides intuition for these surprising dynamic results by showing why their counterparts occur in any given quarter of the model.

[21] We verified that this calibration leads to reasonable investment cost estimates. For instance, our baseline simulations indicate that the total investment cost to achieve the maximum viable capacity (R^*) under the cap is \$24.74 billion. In comparison, Craig Kennedy's estimates indicate that stand-alone export capabilities would cost Russia over \$25 billion. See `https://navigatingrussia.substack.com/p/the-shadow-fleet-in-crisis-highlights` for more details.

[22] We 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.

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 to rely on its shadow fleet.

To gain intuition about the consequences of the two policies, suppose the initial fleet capacity 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_0) \approx 150$) but since Russia would have had no way to export, its producer surplus (revenue net of production costs) would have been zero and, given its investment costs, its profit (producer surplus *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 marginal 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, producer surplus 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 billion per quarter under the ban and $8.86 billion per quarter under the cap) but by the 2^{nd} quarter, Russian profits are higher ($11.72 billion per quarter under the ban and $10.56 billion per 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 profits in each quarter if neither of the policies

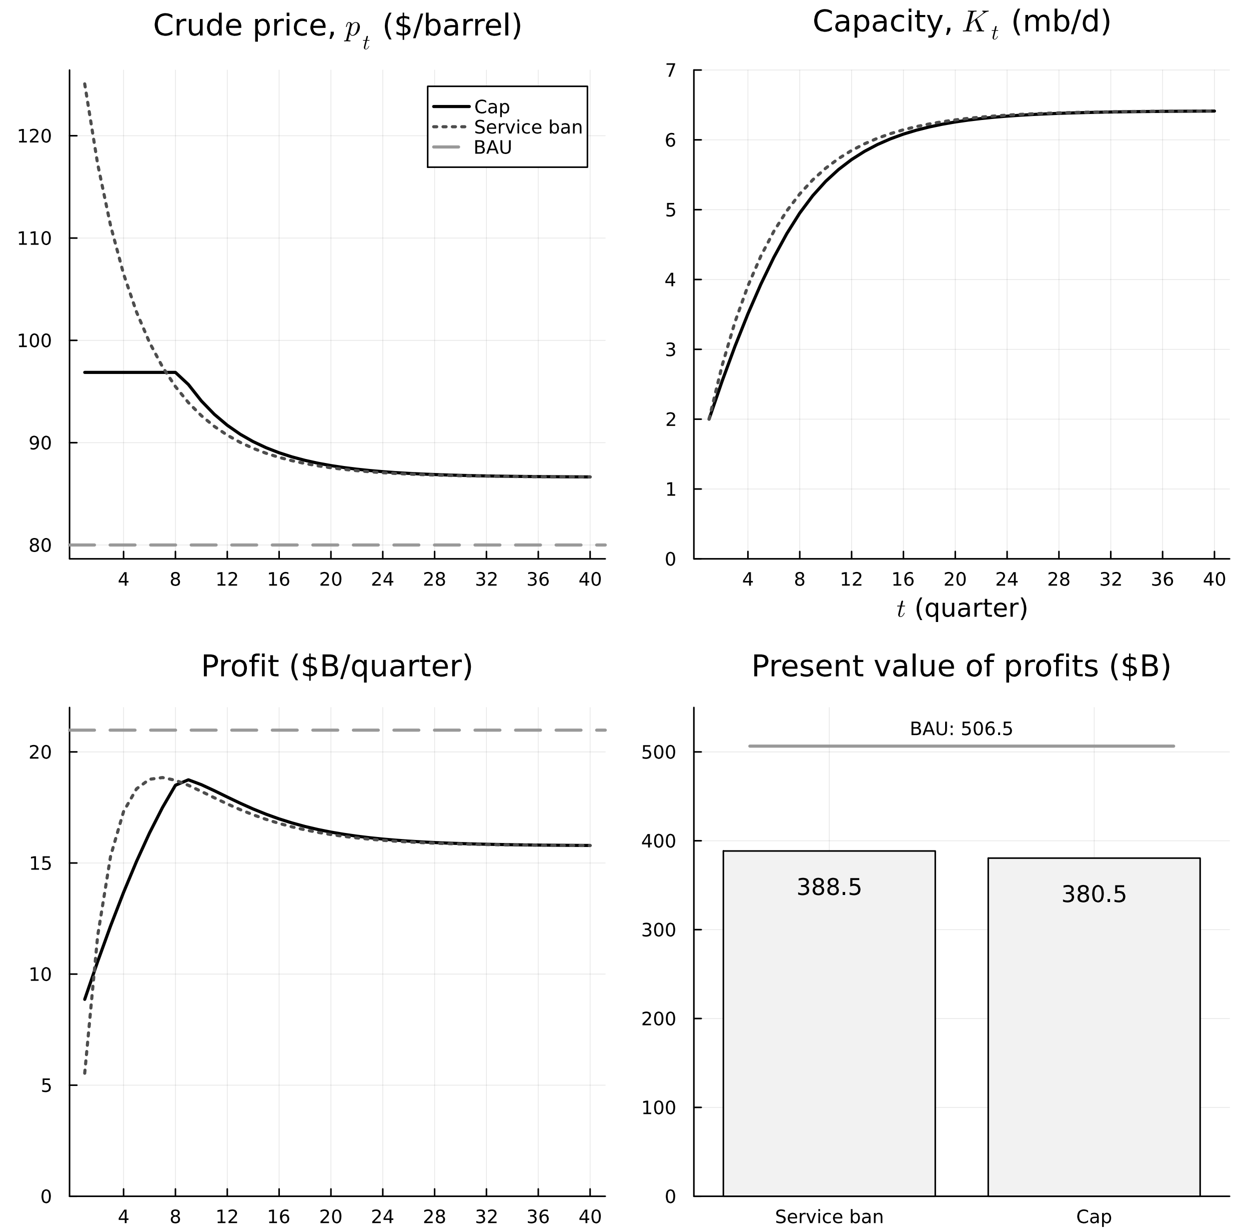

Figure 2: A comparison of prices, capacity, and profits (producer surplus net of investment costs) under a price cap sanction (solid lines) and a service ban (dotted lines). The line graphs display only the first 40 quarters of the 80-quarter simulation.

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 $P(Z_0 + R_c)$ for nearly eight quarters, beginning when the sanction is imposed. These results are robust to changes in the discount rate and the slope of the marginal cost of fleet expansion—see Appendix A for details. However, as Section 6 discusses, the present value of profits under the service ban could become slightly lower than under the cap if the baseline assumptions are relaxed sufficiently.

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 1^{st} 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.

Hence, a cap lower than \$60 would hurt Russia less than a \$60 cap. Indeed, as our simulations with a \$70 cap suggest, a cap slightly higher than \$60 would punish Russia more.[23] In general, a higher cap delays the time when Russia completely abandons its sales at the cap in favor of fleet sales, and 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 price higher than the cap but not $p_t - d$, 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.

[23] Under our baseline assumptions, the 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 1^{st} quarter is monotonic in the cap level.

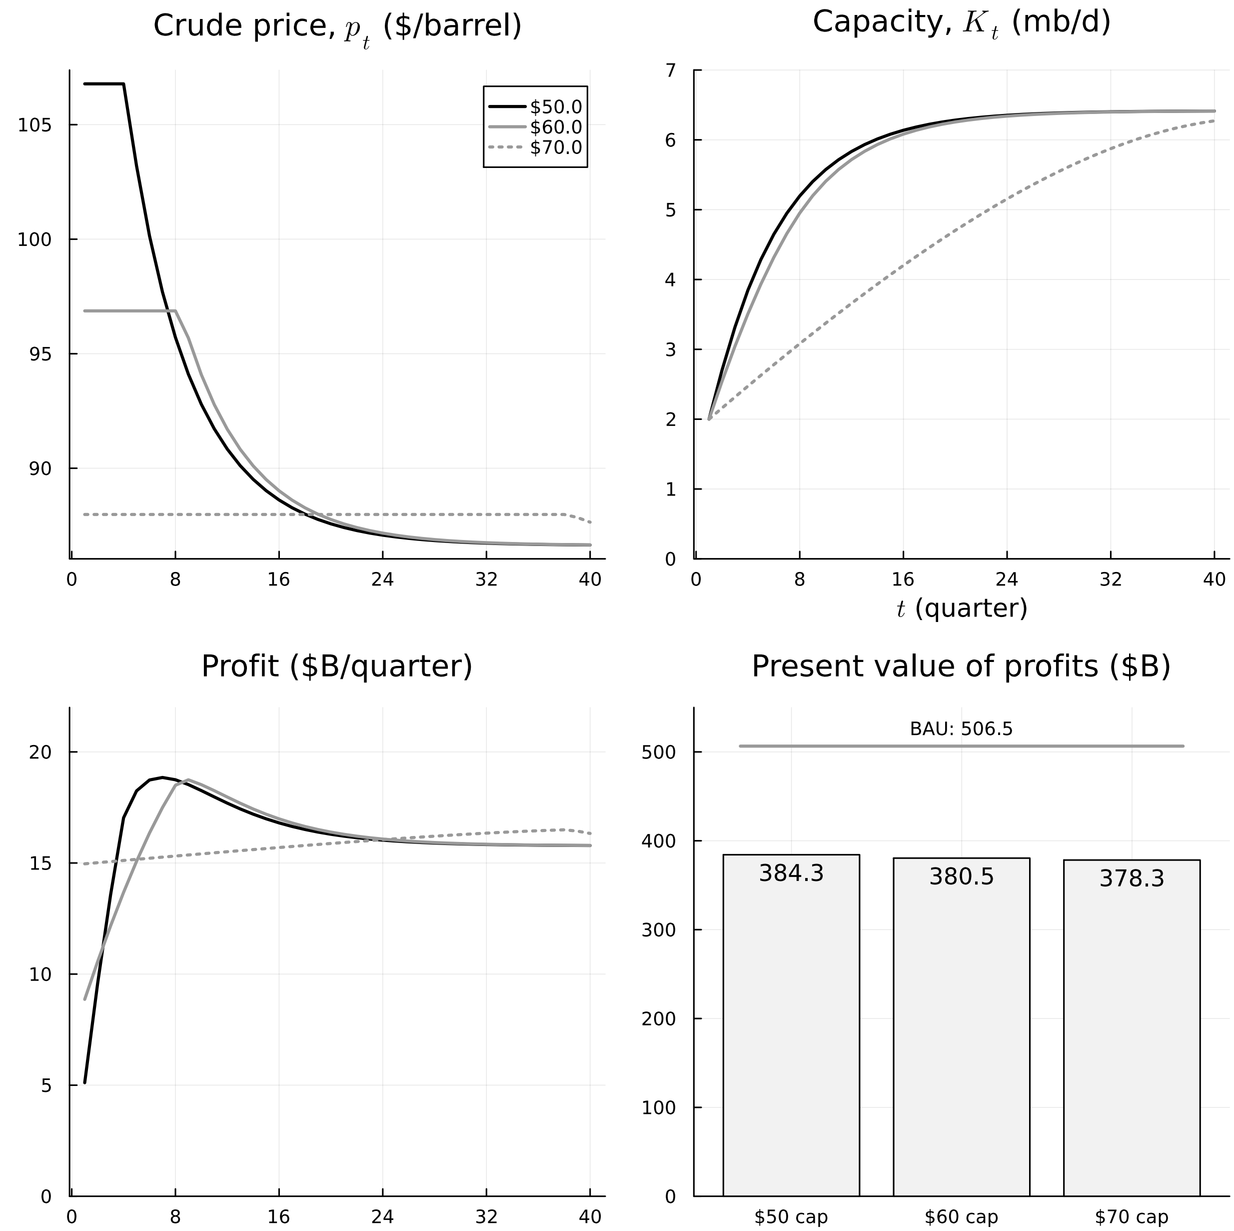

Figure 3: A comparison of prices, capacity, and profits under different cap levels. The line graphs display only the first 40 quarters of the 80-quarter simulation.

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

[24]In practice, it is up to the country conducting the audit to determine the form of the punishment. Some use

at the cap price ($\hat{p}$). Given that cheating is risky, there is a demand for insurance to protect against this risk. One interpretation of d is as a premium on such insurance.

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).

In period t, Russia maximizes

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

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

$$R_t = \quad Q_t + Y_t + X_t, \tag{8}$$

$$X_t \geq 0, \quad (p_t - d) - C'(R_t) - \alpha_t \leq 0, \ c.s., \tag{9}$$

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

$$Y_t \geq 0, \quad [a(\hat{p} - \tau) + (1 - a)(P(Z_0 + R_t) - d)] - C'(R_t) \leq 0, \ c.s.. \tag{11}$$

The conditions defining I_t and α_t are identical to those in the baseline model (equations (3) and (4)).

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 $[a(\hat{p} - \tau) + (1 - a)(P(Z_0 + R_t) - d)] = C'(R_t)$. 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 $\hat{p} = a(\hat{p} - \tau) + (1 - a)P(Z_0 + \hat{R})$. 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 $a(\hat{p} - \tau) + (1 - a)P(Z_0 + Y_t + K_t) = C'(Y_t + K_t)$. 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

fines, others use another sanction.

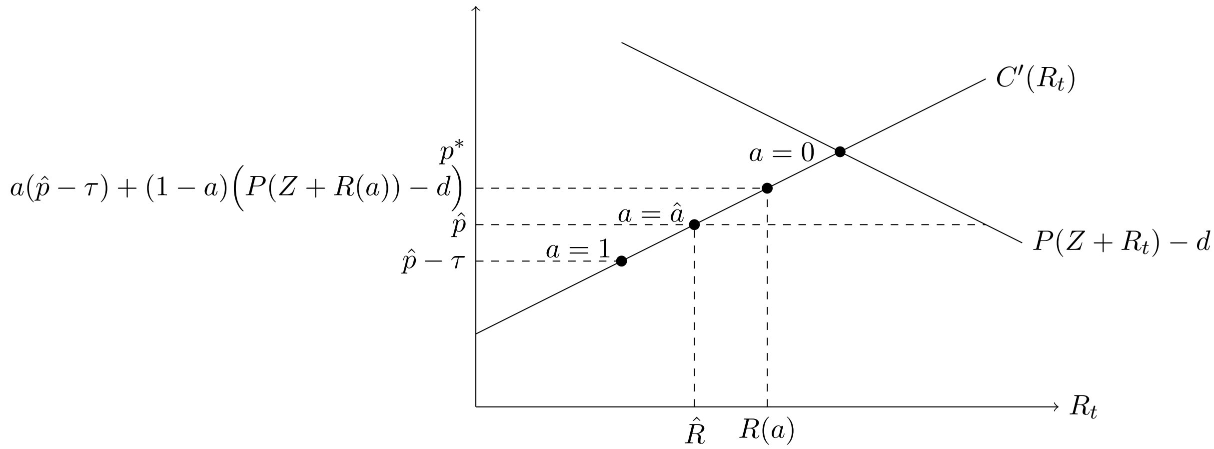

Figure 4: Audit probabilities lower than $\hat{a}$ make cheating strictly more profitable than compliance with the cap; as the audit probability is lowered further, the expected profit increases and more cheating occurs.

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 1^{st} quarter.[25] The expected

[25] While it may be difficult to see in the figures for large t that the fleet capacities never cross, a continuous-time

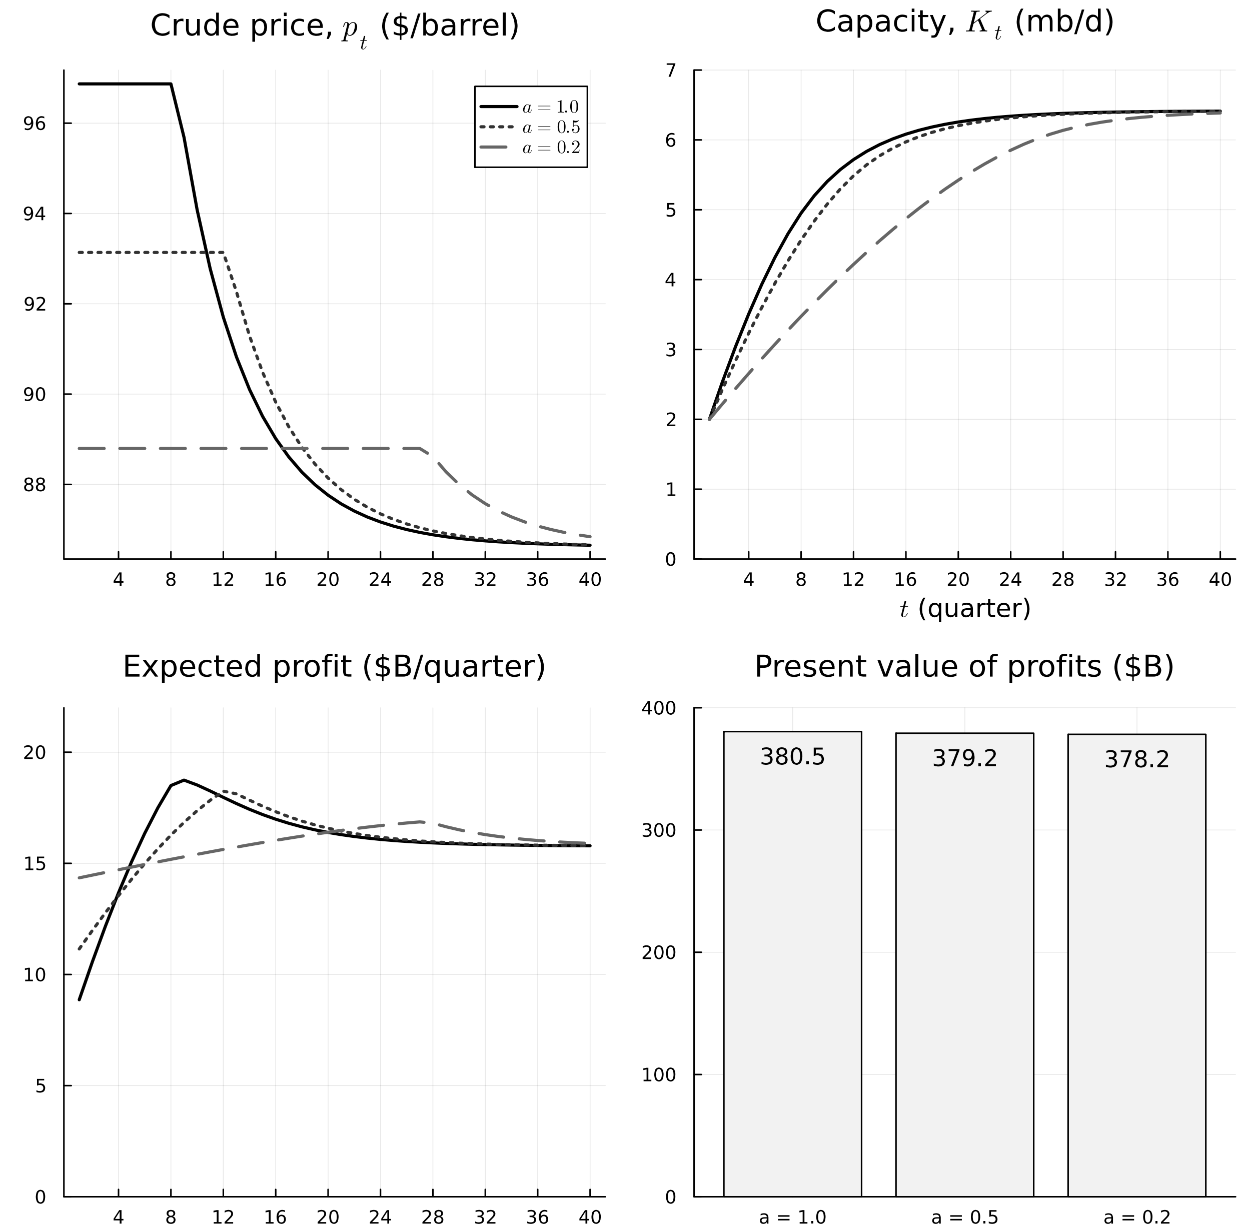

Figure 5: A comparison of outcomes under different levels of enforcement of the $60 price cap sanction. Parameter a is the probability of an audit. All scenarios assume a price cap sanction and a monetary fee of $\tau = \$10$ per barrel. The line graphs display only the first 40 quarters of the 80-quarter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

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

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.

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 5^{th} quarter and the 19^{th} quarter (more than three years), Russian profits are actually higher if enforcement is stringent than if it is lax. So, surprisingly, lax enforcement hurts Russia slightly more in terms of the present value of its expected profits.

4.4 Static Intuition for the Surprising Dynamic Result

In this subsection, we deduce a necessary and sufficient condition for tightening the ceiling to raise static producer surplus in any quarter t, given the fleet size inherited in that quarter. To illustrate, consider how tightening the cap affects static producer surplus in the 1^{st} quarter.

When the cap is tightened, 1^{st}-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 1^{st}-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 producer surplus in the 1^{st} quarter if and only if:[26]

$$\frac{Q_1}{\bar{k}} < -\frac{P'(Q_1 + \bar{k} + Z_0)}{\gamma}. \tag{12}$$

This is not merely a local result. Suppose inequality (12) holds in the 1^{st} 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 (our assumption) and is unchanged if the inverse demand curve was instead linear. So the strict inequality (12) will continue to hold for any lower price ceiling.

That being said, inequality (12) depends on the model's parameters. We verify that it holds for our baseline analysis.[27] However, condition (12) may not hold for other parameters or functional

[26] Since $\hat{p} = C'(Q_1 + \bar{k})$, $\frac{dQ_1}{d\hat{p}} = \frac{1}{C''}$. In equilibrium, 1^{st}-quarter profits are: $\pi_1 = \hat{p}Q_1 + \bar{k}[P(Q_1 + \bar{k} + Z_0) - d] - C(Q_1 + \bar{k})$. Differentiating, we conclude: $\frac{d\pi_1}{d\hat{p}} = [\hat{p} - C'(Q_1 + \bar{k}) + \bar{k}P'(Q_1 + \bar{k} + Z_0)]\frac{dQ_1}{d\bar{k}} + Q_1$. 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 $C'' = \gamma$, $\frac{d\pi_1}{d\hat{p}} = Q_1 + \bar{k}\frac{P'(Q_1+\bar{k}+Z_0)}{\gamma}$. 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 $\frac{d\pi_1}{d\hat{p}} < 0$.

[27] Note that this result is not apparent in Figures 2 and 3 because the increase in the 1^{st}-quarter producer surplus is more than offset by the higher 1^{st}-quarter investment costs induced by the tighter cap. Hence, as these figures

forms. 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, tightening the cap would reduce static producer surplus.[28] In terms of Figure 6 area B would exceed area A.

These static results suggest why a tighter cap harms Russia less in our baseline simulations as well as under a large set of parameters and extensions. They also motivate our dynamic analysis: Indeed, besides its static effect in every quarter, given the inherited evasion capacity, a tighter cap prompts a more rapid expansion of the fleet, affecting Russia both because of the associated investment costs and because this expansion changes the future size of the fleet.[29] Finally, these static results suggest that if our baseline assumptions were relaxed, static producer surplus and, therefore, the present value of Russian profits might *fall* when the cap is tightened. We verify in Section 6 that this conjecture is true, but we find that the quantitative effects of plausibly relaxing our baseline assumptions do not significantly affect our main result about the quantitative importance of the sanctions.

In short, since tightening the cap generates opposing effects, either the cap or the ban may harm Russia slightly more. Given this finding, research time would be better spent devising direct sanctions on the shadow fleet instead of debating the optimal height of the cap.

5 Sanctioning the Shadow Fleet

In this section, we use our simulation model to consider policies that affect the shadow fleet directly.

5.1 The Value of Being Able to Amass a Shadow Fleet

We begin by assessing the benefit Russia derived from being able to expand its shadow fleet after the invasion. Suppose Russia lacked this ability and had to operate with its initial shadow fleet. In that case, there would be no changes in *any* variable over time since the dynamics are driven

reflect, a tighter cap reduces 1^{st}-quarter profits (producer surplus *minus* investment costs).

[28] In footnote 37, we generalize inequality (12) and show that tightening the cap might instead reduce static producer surplus if non-Russian supply is sufficiently price-sensitive ($Z' > 0$) or if the discount importers secure by bargaining is sufficiently sensitive to the gap between the world price and the ceiling price ($d'(p_1 - \hat{p}) > 0$).

[29] In a previous draft, we said that if the ban had a higher producer surplus than the cap in the 1^{st} quarter then, for any γ, the ban must have a higher present value than the cap. This is mistaken, and we thank an anonymous referee as well as G.C. van der Meijden for pointing this out. Instead, it is still possible for the cap to have a higher present value than the ban for some $\gamma > 0$. Whenever this occurs, the present value of the cap as a function of γ must cross the present value of the ban as a function of γ from below when γ is low and from above when γ is higher. In any case, as $\gamma \to \infty$, the ban must again have a higher present value than the cap.

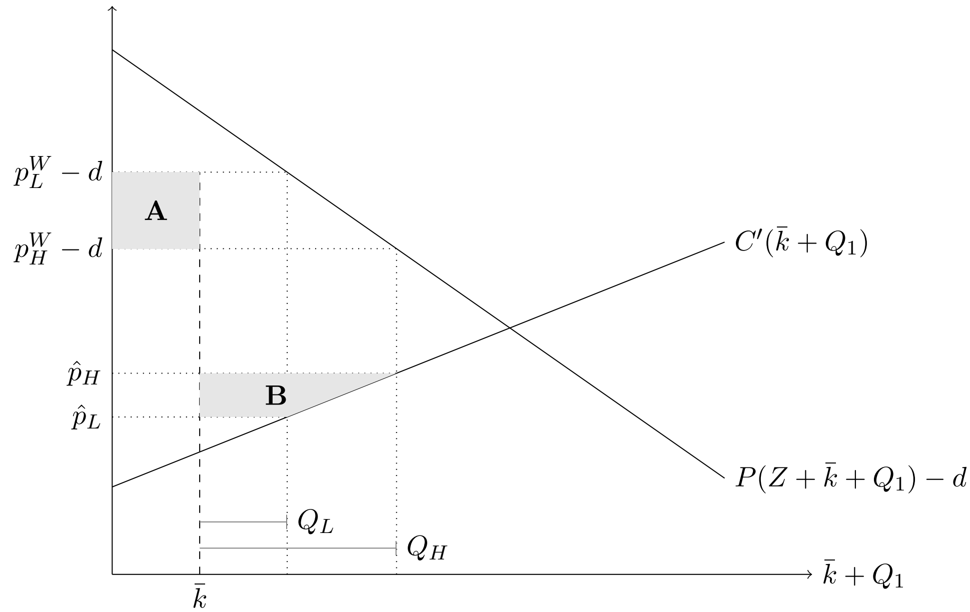

Figure 6: Tightening the ceiling from $\hat{p}_H$ to $\hat{p}_L$ causes 1^{st}-quarter sales using Western services to contract from Q_H to Q_L driving up the world price by $p_L^W - p_H^W > 0$. Hence, producer surplus on fleet sales using the inherited capacity increase by area A but producer surplus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

exclusively by the expansion of the shadow fleet. If the initial fleet size was frozen at $K_1 = 2$ mb/d, the present value of Russian profits over the 80 quarters would have been \$331 billion instead of the baseline \$380.5 billion, a reduction of 13%.

This conclusion rests on our assumption that the fleet capacity was initially 2 mb/d or 27% of Russia's pre-invasion total exports.[30] In fact, this assumption about Russia's initial fleet capacity is extremely conservative. This is because our baseline analysis examines the effect of sanctions implemented at the time of the Ukraine invasion in February 2022. Fleet capacity was much smaller then than when sanctions were implemented three quarters later. In Subsection 6.5 we examine the effects of delayed sanctions and of allowing Russia to expand its fleet in anticipation of the sanctions. We find that Russia would expand its shadow fleet by approximately 1.4 mb/d. This accords closely with IEA estimates (International Energy Agency, 2024) as well as those of *The Financial Times*.[31] If our assumption about initial fleet capacity understates its true size, two of

[30]That is, $\frac{\bar{k}}{\bar{k}+Q_0} = 27\%$.

[31]According to *The Financial Times* Russia had amassed approximately 100 tankers prior to the imposition of the sanctions in December 2022, roughly 42% of what would have been needed to maintain its pre-sanction exports. See, *The Financial Times*, December 2, 2022, available at `https://www.ft.com/content/cdef936b-852e-43d8-ae55-33bcbbb82eb6`. We thank an anonymous referee for directing us to this article.

our main conclusions are re-enforced: (1) Tighter caps hurt Russia less than looser ones; and (2) preventing larger initial fleets from expanding results in a smaller reduction in the present value of Russian profits.[32]

5.2 Raising Russia's Marginal Cost of Expanding the Shadow Fleet

Given our assumption that Russia has a fleet capacity of $K_1 = 2$ by the time when sanctions are applied, we examine policies that have been considered to raise the marginal cost of fleet expansion. The larger the marginal cost of expanding the fleet, the more Russia would substitute sales at the \$60 ceiling for more lucrative fleet sales, and hence the lower would be the present value of Russia's profits. The simulations reported in Figure 7, however, indicate that the *quantitative* effects are very modest. Even a ten-fold increase in ϕ would lower the present value of profits by only 2.8%, consistent with the small benefit to Russia of being able to expand the shadow fleet established previously. As Figure 7 reflects, the \$60 cap harms Russia more than the service ban for any ϕ (provided $K_1 = 2$). The dashed curve increases in ϕ because Russia has to rely on fleet sales, which fetch higher prices. The initial price spike from a service ban is the same regardless of ϕ. However, the larger ϕ is, the longer it takes this initial price spike to dissipate.[33]

5.3 Unanticipated Targeting of the Shadow Fleet

Policies have been considered to reduce the size of Russia's shadow fleet without warning. 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.[34] It is useful to divide targeting cases into three categories: (*i*) when Western services are being utilized until the time targeting is implemented; (*ii*) when only the shadow fleet is utilized even after targeting is implemented; and (*iii*) when only the shadow fleet is utilized until targeting is implemented but tankers using Western services are added immediately after targeting is implemented.

In the four panels of Figure 8, the black paths ("no targeting") describe the baseline simulation

[32] Given our other baseline parameters, an initial fleet size of less than 1.57 mb/d at the time of the invasion is necessary for a tighter cap to be more harmful to Russia.

[33] Figure A4 in Appendix A shows the simulated trajectories for different values of ϕ.

[34] Targeting may be paired with changes in the enforcement policy. To save space, we relegate simulations of such policy combinations to Appendix C.

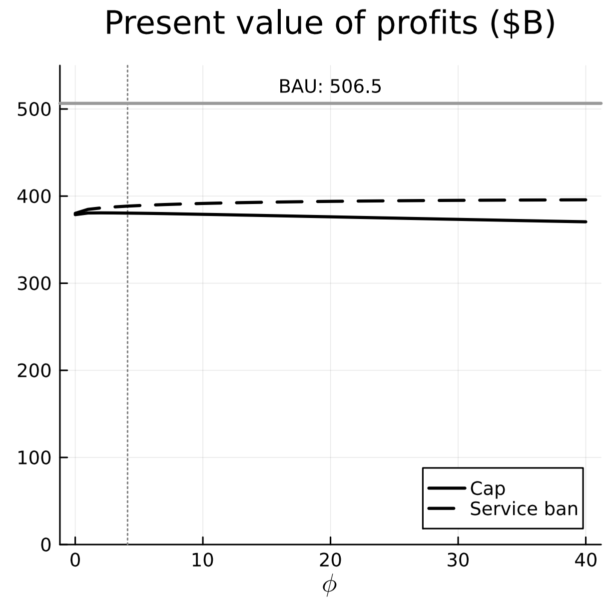

Figure 7: Present value of profits under the two policies as ϕ increases.

and the dashed paths show the effects of targeting initiated at the end of the first phase ($t = 8$) in the baseline simulation ($K_8 \approx 4.95$). Suppose that the fleet size is suddenly reduced to 3.5 mb/d, the level attained in the 4^{th} quarter of the baseline simulation. As the top-left panel reflects, the world price would be unaffected. Producer surplus in the 8^{th} quarter must be lower since production costs are the same, but some exports previously sold at the net world price are sold instead at the ceiling. The present value of Russian profits over all 80 quarters is strictly lower than in the baseline. As the top-right panel confirms, after the downward jump in capacity initiated in the first phase, investment to rebuild the fleet is intense, implying that adjustment costs jump up. Since the producer surplus jumps down and the investment costs jump up, profits (producer surplus net of these adjustment costs) must jump down as the bottom-left panel illustrates.

At the other extreme (case *ii*), suppose that after the unanticipated reduction in capacity occurs, the world price path is still declining (*second* phase). Then the price must jump up. In Figure 8, the reduction occurs in the 12^{th} quarter and the stock is reduced to 4.95 mb/d, which had been attained in the 8^{th} quarter. In this case, the producer surplus in the "initial" period ($t = 12$) increases because the reduction in production costs exceeds the reduction in revenue from fleet sales. Since the stock is closer to its steady-state level, investment in the 12^{th} quarter is less intense than when the targeting occurs earlier and the fleet is smaller. Over all 80 quarters, the

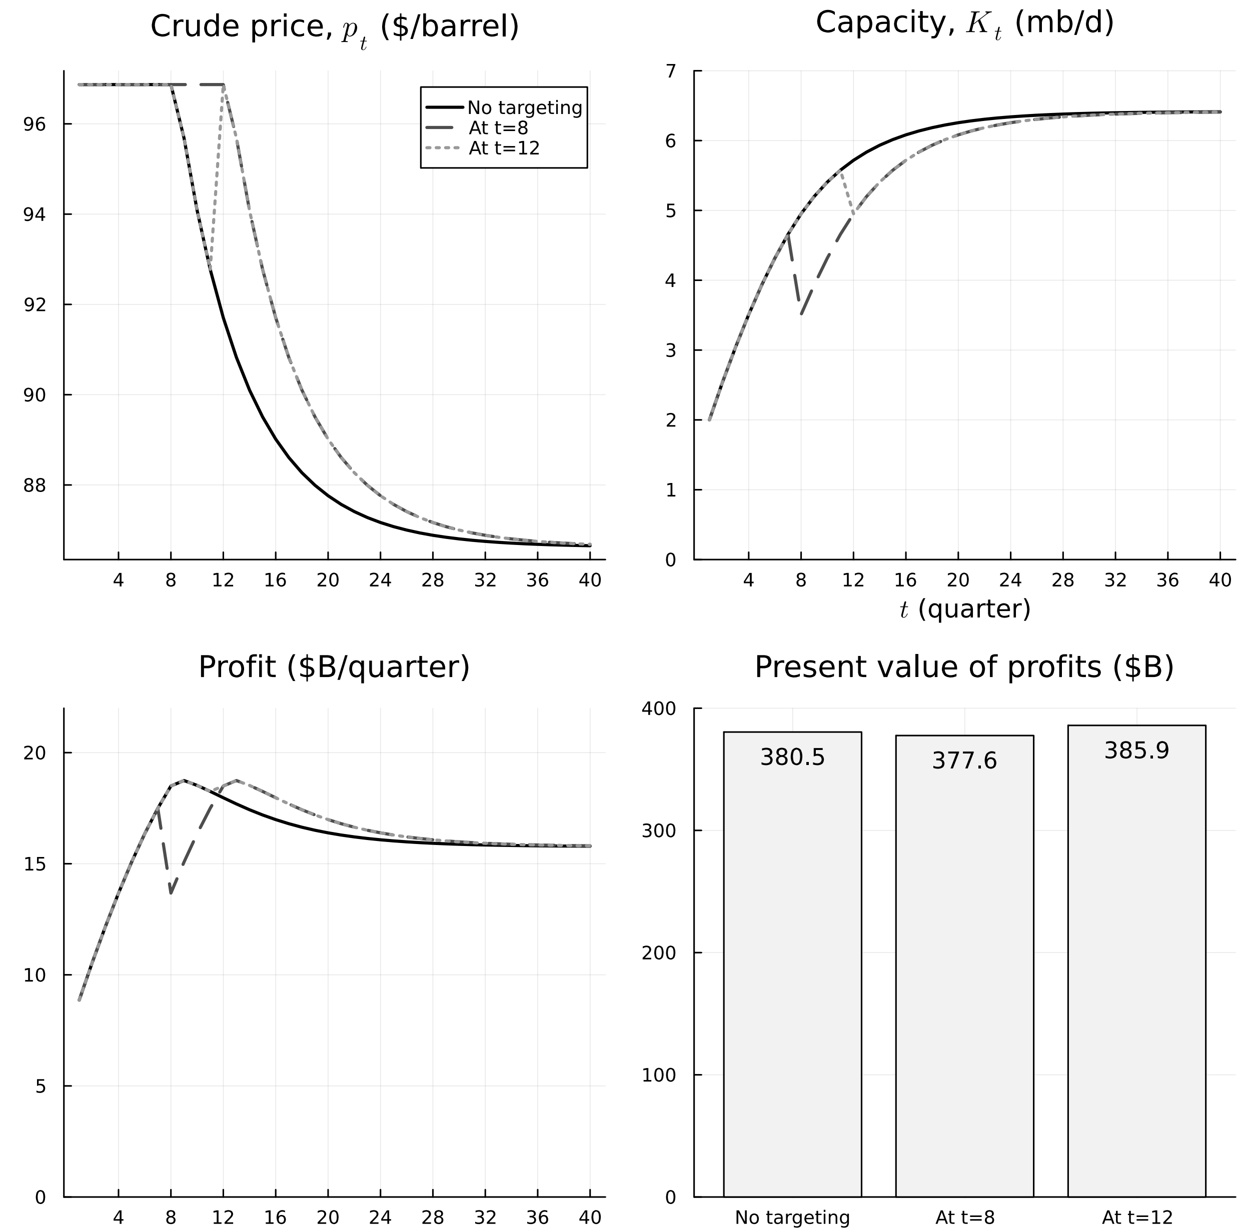

Figure 8: The consequences of unanticipated targeting of the shadow fleet when it begins in the first phase or ends in the second phase.

present value is higher with targeting than in the baseline. This seems counterintuitive since Russia could always have destroyed part of its own shadow fleet. But as before, that strategy would have lowered its discounted profits when assessed at the price path prevailing at the time; that Russia benefits in the new equilibrium from the downward jump in its shadow fleet is a reflection of the change in the price path induced by the change in policy—another manifestation of a surprising equilibrium result.

6 Extensions

In this section, we discuss how our main results would be affected if the parsimonious modeling assumptions in the baseline analysis were modified. In Subsections 6.1–6.5, we discuss the implications of replacing these simplifying assumptions with more flexible functional forms. We also relax our assumption that the sanctions were imposed immediately after the invasion and instead assume that the three-quarter delay before they were actually imposed was perfectly anticipated.

6.1 Non-Russian Supply Response

In the baseline version of our model, we assumed that the non-Russian supply of oil is fixed. In reality, estimates of the price elasticity of global oil supply typically range from 0.01 (Kilian and Murphy, 2014) to 0.15 (Baumeister and Hamilton, 2019). Therefore, we examine here the consequences of assuming an upward-sloping non-Russian oil supply.

In Appendix B.1, we explicitly replace the fixed supply Z_0 by the upward-sloping, isoelastic, supply function $Z(p) : Z = Z_0(p_t/p_0)^{\epsilon_Z}$, where Z_0 and p_0 are defined in the baseline model (see Table 1), and where the new parameter $\epsilon_Z \geq 0$ is the price elasticity of non-Russian supply.

We consider price elasticities of non-Russian supply of $\epsilon_Z = .1$ and $\epsilon_Z = .2$. Compared to the baseline case ($\epsilon_Z = 0$), these elasticities dampen the oil price spikes due to a service ban and, therefore, slow down the shadow fleet expansion. Neither of these elasticity values affects our main result that the sanctions generate significant and similar present-value losses for Russia. However, if non-Russian supply were sufficiently elastic ($\epsilon_Z > 0.105$), the ordering of the policies would be reversed: The service ban would harm Russia slightly more than the price cap.

6.2 Curvature of Marginal Production Costs

Our baseline analysis assumes a quadratic cost of production. This simplifying assumption is well-suited to our calibration exercise: We have calibrated the increasing marginal cost by linearly interpolating the lowest marginal cost obtained from reserve data and the marginal cost inferred from actual, pre-invasion, production data.

Using field-level data, however, Wachtmeister et al. (2023) establish that the marginal cost of Russia's potential production exhibits convexity, with curvature especially pronounced in the short run. The actual form of the marginal cost curve suggests that in our baseline model, supply might be more responsive to price changes near \$80 than it is in reality, with implications for changes in quantity, price, and thus profits.

In Appendix B.2, we replace marginal cost C' by the convex function:

$$\tilde{C}'(X_t + Q_t) = c_0 + (p_0 - c_0)\left(\frac{X_t + Q_t}{Q_0}\right)^{\chi},\ \chi \geq 1, \tag{13}$$

where c_0, p_0, and Q_0 are as defined in the baseline model (see Table 1), and where the new parameter χ introduces a degree of curvature of the marginal cost function in between the two end points used for calibration. The case $\chi = 1$ corresponds to our baseline simplification in which marginal cost is a linear function. Higher values $\chi > 1$ introduce convexity, capturing both easier production adjustments at low production levels and rigidities at high production levels. As χ increases, the marginal cost curve tends towards a "hockey stick" form, as illustrated in Figure 9.

Simulations of our model with $\chi = 2$ or $\chi = 8$ and the other baseline parameters unchanged show that a higher degree of curvature of the marginal costs induces lower world prices and greater profits. The effects on prices can be explained as follows. The introduction of the cap is followed by a smaller reduction in exports subject to the cap, affecting the world price less. Profits are higher because the convexity of marginal cost between the two calibration end points implies a lower total cost of production.

Appendix B.2 shows how various degrees of curvature of the marginal cost reflect the arc-elasticity of Russian oil supply between the end of 2021 (pre-invasion) and the end of 2023. Comparison with the supply arc-elasticity based on actual produced quantities and Urals prices suggests that our linear simplification of the marginal cost curve overestimates the Russian supply response.

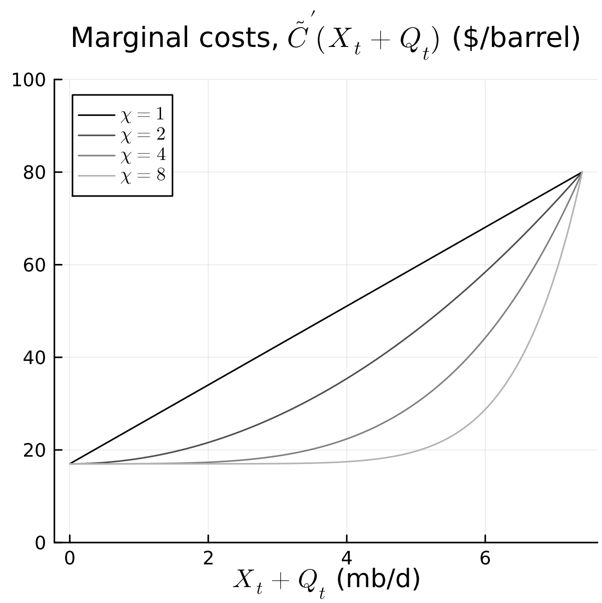

Figure 9: Varying curvature of marginal production costs with different levels of parameter χ.

A curvature parameter of $\chi = 2$ generates a supply response close to what was observed.

Nonetheless, our main result—that both the $60 price cap and the service ban impact Russia's present value of profits significantly and in a similar ways—persists despite the increased curvature of the marginal cost. A higher degree of curvature is required for the ban to harm the present value of Russian profits more than the price cap, although the difference in profits remains small.

6.3 Endogeneity of the Sales Discount

Recall that Russia receives d less than the world price on its shadow fleet exports and, in the case of imperfect enforcement, also receives d less than the world price on barrels shipped in violation of the cap. One interpretation of d in the imperfect enforcement case is as a premium on insurance against being apprehended while cheating.

In our baseline simulations, the discount at which Russia sells its exports via the shadow fleet is held fixed at $d = 15$, reflecting the differential between Brent and Urals reference prices when it stabilized approximately one year after the actual start of sanctions. In addition to an insurance premium, the literature suggests this discount could also reflect a combination of factors, including higher shipment costs and monopsony power.[35] Kilian et al. (2024a) estimate the increased cost of

[35] Buyers in India and China, purchase the limited quantities of Russian oil available to them at $p_t - d$ but must

shipping to more distant buyers, finding higher costs notably to China and India. A higher discount may also indicate the increased cost of insurance against being found in violation of the cap.

Johnson et al. (2023) suggest that this discount partly results from the bargaining power of non-coalition importers of Russian oil, observing that the price cap "likely provides negotiating power to oil importers that continue to buy Russian oil above the cap without using [Western] services." Kilian et al. (2024a) substantiate that the bargaining positions accounted for observed high initial discounts on Russian oil. Arguably, the opportunity cost of Russia having to sell at the capped price if negotiations with importing countries broke down might account for this bargaining power.

Suppose we assume that the discount is an increasing function of the difference between the world price and the cap. Then if the discount reflects bargaining, the lower the cap, the larger the discount from bargaining. In particular, in the first phase, a reduction in the cap would cause Russia to reduce its sales at the cap, driving up the world price and hence driving up d. In the second phase, nothing is sold at the cap, and tightening it has no effect. Suppose instead the discount reflects the premium on insurance in the imperfect enforcement case. If the audit probability a were increased, the expected price per barrel would fall, less cheating would occur, and the world price would rise. Hence, the discount Russia receives "even on the non-shadow fleet" would rise "due to the risk of being found in violation of the cap."[36]

Appendix B.3 examines how our results are affected by the sensitivity of the discount on purchases shipped by the shadow fleet. To calibrate this sensitivity, we assume that the discount d is an affine function of the difference between the world price and the lower price alternative ($60 for cap scenarios and marginal production cost for ban scenarios). We calibrate the function $d(p-\hat{p}) \equiv \nu_0 + \nu\,[P(Z_0 + K_t) - \hat{p}]$, exploiting variations in the temporary differential between Brent and Urals reference prices and attributing the stabilized discount to factors other than bargaining power. Our calibration yields $\nu = .33$ and $\nu_0 = 6.32$. This functional form implies that when $Q_t = 0$ and $X_t = R^* = 6.4$ mb/d, the world price is $P(Z_0 + R^*) = \$86.62$. The discount is then equal to its baseline value of $d = \$15$.

pay the world price p_t to satisfy their remaining demand. Hence, whether the discount d reflects shipping costs, insurance premia, or bargaining, buyers would pay the marginal price p_t and Russia would receive $p_t - d$.

[36]The quotations are from an extremely perceptive referee, whom we thank for pointing out this implication of our model.

We find that this formulation of the discount is consistent with the data. Using this calibration, our simulations confirm that both the price cap and the service ban substantially reduce Russian profits. A more sensitive d is required for the service ban to harm Russia's present value of profits more than the price cap, although the difference in profits remains small.[37]

6.4 Internalization of Price Effects

The baseline analysis assumes that Russia acts as a price taker. Despite its stated commitment towards OPEC+, Russia's oil supply behavior has been more consistent with our baseline price-taking hypothesis than with the exercise of pure monopoly power—see, e.g., Cahill (2023). This is partly explained by the hybrid structure of the oil industry in which oil supply decisions are largely decentralized, leading to coordination issues (Henderson and Fattouh, 2016).

Yet, Russia's market share suggests that it has some market influence. We discuss here how our results would change had we assumed that Russia internalizes its market influence to varying degrees.

We model this partial market influence in Appendix B.4, by extending our model to a mark-up $\mu_t \equiv \theta X_t P'(Q_t + X_t + Z_0)$, proportional to the mark-up that a monopoly would impose ($X_t P'(Q_t + X_t + Z)$), capturing the degree $0 \leq \theta \leq 1$ in which Russia manages to exert its potential market power. In the polar case in which $\theta = 0$, Russia behaves as a price taker and the conclusions of the baseline simulations re-emerge. When $\theta = 1$, Russia fully internalizes its market influence, as if its industry were perfectly coordinated, behaving like a pure monopoly.

We establish that our main result would not be significantly affected by partial market influ-

[37] Had we considered simultaneously price-responsive non-Russian oil supply, an endogenous discount to buyers, and a curved marginal cost, the profit function would be extended to:

$$\pi = \hat{p}Q + K[P(Q + K + Z(p)] - d(p - \hat{p})] - \tilde{C}(Q + K). \tag{14}$$

Differentiation with respect to $\hat{p}$ and use of the first-order condition yields:

$$\frac{d\pi}{d\hat{p}} = \underbrace{\overbrace{Q}^{\text{Hotelling}} + KP'\frac{dQ}{d\hat{p}}}_{\text{Baseline}} + \overbrace{\frac{dP}{d\hat{p}}[KP'Z' - v] + v}^{\text{strictly positive terms}}, \tag{15}$$

where

$$\frac{dQ}{d\hat{p}} = \frac{1}{\tilde{C}''}. \tag{16}$$

If $Z' > 0$, $v > 0$, or $\tilde{C}'' > C''$, equation (15) implies that inequality (12) is less likely to hold.

ence: Both the price cap and the complete service ban reduce the present value of Russian profits substantially. Moreover, the result that a lower cap would increase Russian profits emerges only for high degrees of market power internalization ($\theta > 0.85$).

6.5 Anticipatory Behavior

Finally, in our baseline analysis, we assumed that sanctions were implemented immediately after the February 2022 invasion. In fact, the cap was implemented in December 2022. G-7 discussions and early consultation with market participants (Wolfram, 2024) made everyone keenly aware of the coming sanctions—including Russia, which used the nine-month delay as an opportunity to continue selling its seaborne exports at the world price while simultaneously expanding its shadow fleet.

In Appendix B.5, we use our model, calibrated with its baseline parameters, to simulate the consequences of anticipating this policy of delayed sanctioning. That is, we assume the policy Russia anticipated is to be sanction-free during the first three quarters and then to have a $60 price cap (or, alternatively, a service ban) imposed by the West from the 4^{th} quarter onward.

Since the fleet would not have been used during the first three quarters, Russia would choose its fleet expansion program to minimize the present value cost of achieving a fleet of a given size by the end of the 3^{rd} quarter. This is why the marginal cost of fleet expansion ($F'(I)$) increases during the first three quarters by a constant rate of 4.15% (the reciprocal of the quarterly discount factor). The initial investment (I_1) depends on the policy. For example, it is larger when the delayed policy is a service ban rather than a cap since a future ban would leave Russia reliant entirely on its shadow fleet for its seaborne exports. Once sanctions are imposed and the shadow fleet is utilized, the marginal cost of expansion begins to decrease. In terms of Table B2 and the top right panel of Figure B10, fleet capacity increases at an increasing rate for the first three quarters and then increases at a decreasing rate. By the time the anticipated sanction is imposed, the fleet would have grown from 2 md/d to 3.38 mb/d in the case of a $60 cap and to 3.65 mb/d in the case of a service ban.

In the bottom left panel of Figure B10, quarterly profits decline during 2022 because producer surplus remains at the *laissez-faire* level while the cost of investment rises at an increasing rate in successive quarters. Delaying the sanction by three quarters raises the present value of Russia's profits by \$19.2 billion in the case of the \$60 cap and \$14.7 billion in the case of a service ban.

In the absence of sanctions, Russia would have earned \$506.5 billion (the BAU level) in the present value of profits. With a \$60 cap, the harm to Russia's present value of profits very much depends on when the cap was imposed and whether Russia was prevented from using or expanding its shadow fleet.

Suppose the cap had been imposed immediately following the invasion. If Russia had been prevented from using its shadow fleet, the present value would have fallen to \$236 billion, a reduction of \$270.5 billion. Russia could have restored \$95 billion (35% of that reduction) by using the shadow fleet capacity of 2 mb/d available even if Russia had been prevented from expanding it. If Russia could have expanded that capacity, then Russia would have restored an additional \$49.5 billion (18.3% of the original reduction).

In fact, the \$60 cap was not imposed immediately but three quarters after the invasion. To the extent that imposition of the cap then was perfectly foreseen, Russia was able to recover an additional \$19.2 billion (7% of the original reduction). In total, therefore, Russia recovered 60.3% of the potential loss in present value that could in principle have been achieved with the \$60 cap.[38]

Appendix B.5 also reports on the effects of a service ban delayed until nine months after the invasion. It harms Russia slightly less than a delayed cap. Russia's present value was cut to \$403.2 billion (by only 20.4% of the BAU value). Thus, as in our simulations without delays, the delayed ban harms Russia slightly less than the delayed cap, but the difference pales in comparison to the harm each policy does when compared to BAU.

In our simulations of the delayed cap, the price jumps from the *laissez-faire* level of \$80 per barrel in $t = 3$ to the price plateau of \$96.87 one quarter later, a 21% increase. This jump is heavily dependent on our assumption that no agent anywhere in the world finds it feasible to buy oil in the quarter before the cap is imposed, store it aboveground, and then sell it one quarter later.[39] In

[38] If the cap was imposed at the time of the invasion, Russia's present value would have been \$331 billion if capacity was frozen at 2 mb/d and \$380.5 billion if the capacity could have been expanded from that initial level. Since the cap was imposed after a three-quarter delay, Russia's present value was \$399.7 billion.

[39] If agents somewhere in the world could make a strict profit by this intertemporal arbitrage, then there would be positive storage in equilibrium. Accumulation of inventory would begin prior to the imposition of sanctions, and decumulation of inventories would begin (and possibly end) on the quarter when sanctions were imposed ($t = 4$). The number of quarters during which inventories would be held equals the number of periods required for the price to rise 21% (from \$80 to \$96.87) at the constant percentage rate required to cover interest and (iceberg) storage costs. If, for example, a 10% capital gain per quarter was necessary, storage would have occurred over two quarters. The cap is imposed on the 4^{th} quarter. In that quarter, decumulation of inventories would be necessary to replace the sharp reduction in Russian exports that would occur when the marginal cost of production jumps down from the \$80

our simulations, we made this assumption in order to isolate the effects of shadow fleet expansion since that phenomenon is both novel and significant whereas the dynamic effects of storage are well-known and likely of little quantitative significance.

7 Conclusion

Since invading Ukraine in February 2022, Russia has used profits from its seaborne oil exports to finance the war. In December 2022, the West responded by boycotting Russian seaborne oil exports and by refusing to allow Russia to use Western services (e.g., insurance) on tankers carrying Russian oil selling for more than \$60 per barrel. Since this so-called "price cap" is without historical precedent, economists have relied on a combination of theory and simulation to understand its effects.

We have built and calibrated a tractable simulation model to assess the effectiveness of the price cap policy and other policies intended to reduce Russia's ability to finance the war. The centerpiece of our dynamic model is Russia's ability to expand its "shadow fleet"—tankers that use no Western services and are therefore not subject to the price cap. Previous literature had implicitly assumed that Russia lacked the ability to expand its shadow fleet.

Among the policies we considered are price caps set at various levels, a service ban (the equivalent of a price cap set sufficiently low), alternative auditing intensities to enforce the cap, and sanctions directly targeting the use and expansion of the shadow fleet. Our model can be used in two modes: immediate imposition of sanctions and a foreseen delay in the imposition of sanctions. By assuming immediate imposition of sanctions and an initial evasion capacity of its presumed size in February 2022, we were able to simulate what would have happened if sanctions had been imposed immediately following the invasion. By assuming an anticipated delay of three quarters of a year and the same initial fleet capacity of 2 mb/d, we were able to simulate the consequences of the actual policy of a foreseeable delay of the cap until December 2022. We were thus able to compare the effects of different sanctions, whether they were imposed immediately or after a given delay. We were also able to assess the consequences of delaying any given sanction instead of

BAU price to the \$60 cap as per equation (1). If that decumulation exhausted the inventory, the cap would have been imposed when the world price reached \$96.87. If stocks remained, the cap would have been imposed when the world price was lower, and it would continue to rise in subsequent quarters to cover interest and storage costs until it reached the price plateau of \$96.87.

imposing it immediately.

Our primary finding is that the $60 price cap, the service ban (equivalent to a cap below $34), and caps set at levels in between all reduce the present value of Russian profits by more than $100 billion regardless of whether the sanction is imposed immediately or delayed. However, perhaps surprisingly, each sanction reduces the present value by approximately the same amount. What drives this result is the induced effect on the world price.

When the price cap is lowered, Russia sells less at the cap, and its exports fall. While this ensures that the revenue from cap sales falls, there is a countervailing effect that is typically overlooked. Since Russian exports fall when the price cap is lowered, the world price rises, and this raises the revenue from shadow fleet sales. Because these two effects work in opposite directions, the net effect on the present value of Russian profits from a lower cap may be small, and its sign may even depend on parameters.

We can use our calibrated model to ask whether an outright delayed service ban, the policy first proposed by the EU, would have been a better policy choice than the delayed $60 cap that was ultimately adopted. Given our baseline parameters, the delayed $60 cap would have harmed the present value of Russian profits more than a delayed service ban (or its equivalent, $34 cap) by a mere $3.5 billion. A delayed cap lower than $60 would have further diminished this slight advantage.

While a tighter delayed cap would have had little effect on the present value of Russian profits, the price increase it would have induced is significant. If a $34 price cap were imposed, the world price in our baseline simulations would have been 29% higher than after imposition of a $60 cap. The higher price would have harmed consumers by raising the price of gasoline and the price of other consumer products that use oil as an input. Thus, our baseline parameters imply that a lower delayed price cap would have harmed Russia less while at the same time harming consumers more than a higher cap.

The price increase that tightening the delayed cap would have induced, according to our baseline model, may be overstated. This would occur, for example, if our baseline parameters understate the supply elasticity of non-Western producers of oil or the demand elasticity of global demand for oil. In either case, a smaller increase in the world price would have sufficed to eliminate the excess demand arising when Russia contracted its exports in response to a lowered cap or an outright service ban. Thus, it is possible that if our baseline parameters were revised as discussed in the previous section, a lower delayed price cap would have harmed Russia slightly more and consumers less in 2022 than in our baseline simulations.

But it definitely does not follow that a lower cap imposed *today* (or announced *today* and imposed three quarters later) would now serve Western interests. That is because the shadow fleet now is considerably larger than the 2 mb/d used in these historical simulations and, as a result, the harm to Russia from a tighter cap will be smaller. Since the current shadow fleet is relatively large, sales at the ceiling are relatively small. In this circumstance, a reduction in the price cap would lower the revenue from cap sales by a small amount but would raise fleet sales revenue by a large amount.

Although different levels of a delayed cap affect the present value of Russian profits very little, the delay itself benefits Russia by increasing the present value of Russian profits significantly. For example, the postponement of the $60 cap and accompanying embargo during 2022 allowed Russia to export all of its seaborne oil at the world price while at the same time expanding its shadow fleet over nine months by approximately 70%. Had the sanctions instead been imposed immediately, the present value of Russian profits would have been $19.2 billion lower. This is significantly larger than the benefits from even extreme changes in the level of the cap.

This suggests there is a benefit to devoting less time to fine-tuning the height of the cap and more to shortening the delay in imposing new sanctions.

References

Babina, T., Hilgenstock, B., Itskhoki, O., Mironov, M., and 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

Baumeister, C., and Hamilton, J. D. (2019). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review*, *109*(5), 1873–1910.

Becker, R. A. (1985). Capital income taxation and perfect foresight. *Journal of Public Economics*, *26*(2), 147–167.

Brander, J., and Spencer, B. (1985). The role of inventories and speculative trading in the global market for crude oil. *Journal of International Economics*, *18*(1-2), 83–100.

Brooks, R., and 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)

Cahill, B. (2023, June). Opec has a russia problem. *Center for Strategic and International Studies (CSIS).* Retrieved from https://www.csis.org/analysis/opec-has-russia-problem

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., and 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

Henderson, J., and Fattouh, B. (2016, February). Russia and OPEC: Uneasy partners. *Oxford Institute for Energy Studies*. Retrieved from https://www.oxfordenergy.org/wpcms/wp-content/uploads/2016/02/Russia-and-OPEC-Uneasy-Partners.pdf

Hilgenstock, B., Ribakova, E., Shapoval, N., Babina, T., Itskhoki, O., and Mironov, M. (2023).

Russian Oil Exports Under International Sanctions. *SSRN*. Retrieved from https://ssrn.com/abstract=4430053

International Energy Agency. (2022). *Oil Market Report - January 2022.* Paris. Retrieved from https://www.iea.org/reports/oil-market-report-january-2022

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., and Wolfram, C. (2023). 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., and Wolfram, C. (2025). *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

Johnson, S., and Wolfram, C. (2024, May). *Strengthening enforcement of the russian oil price cap* (Tech. Rep.). Brookings Institution. Retrieved from https://www.brookings.edu/wp-content/uploads/2024/05/20240528_ES_Sanctions_JohnsonWolfram_Final.pdf

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., and Murphy, D. P. (2014). The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, *29*(3), 454–478.

Kilian, L., Rapson, D., and 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., and 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)

Salant, W. A. (1942, December 17). The selection of industrial bombing targets: Some analytical notes. In W. W. Rostow (Ed.), *Pre-invasion bombing strategy: General eisenhower's decision of march 25, 1944* (pp. 99–104). Austin: University of Texas Press.

Shatz, H. J., and 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., and 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., and Prescott, E. C. (1989). *Recursive methods in economic dynamics.* Harvard University Press.

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

Wachtmeister, H., Gars, J., and 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]

Contents

[a]Department of Agricultural and Consumer Economics, University of Illinois Urbana-Champaign.

[b]Department of Economics, University of Michigan, and Resources for the Future (RFF).

[c]Department of Technology, Management and Economics, Technical University of Denmark, Center for Energy and Environmental Policy Research, Massachusetts Institute of Technology, and CESifo. Corresponding author: jxada@dtu.dk.

A 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.

A.1 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 price elasticities of demand (ϵ_D) varying in absolute value between 0.05 and 0.75—six times larger than the baseline elasticity of -0.125.[1] The curves in Figure A1 compare the present value of Russian profits under a service ban and a $60 cap for a wide range of elasticities. In this sensitivity analysis, we recalibrate ϕ in each run so as to reproduce the same procedure used in the baseline to match initial and ninth-period capacity estimates.

Based on Figure A1 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 flipping the ranking, with the cap policy delivering slightly larger present value of profits than then service ban.[2] Further inspection shows that these policies result in the same present value of profits at $\epsilon_D \approx 0.23$.

The second key message from Figure A1 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

[1] In the paper, we report that removing all Russian oil from the market would have driven the world price to approximately $150 ($P(Z_0) \approx 150$) under the baseline elasticity. In comparison, an elasticity of -0.05 or -0.75 would lead this price to $376 or $89, respectively.

[2] For 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.

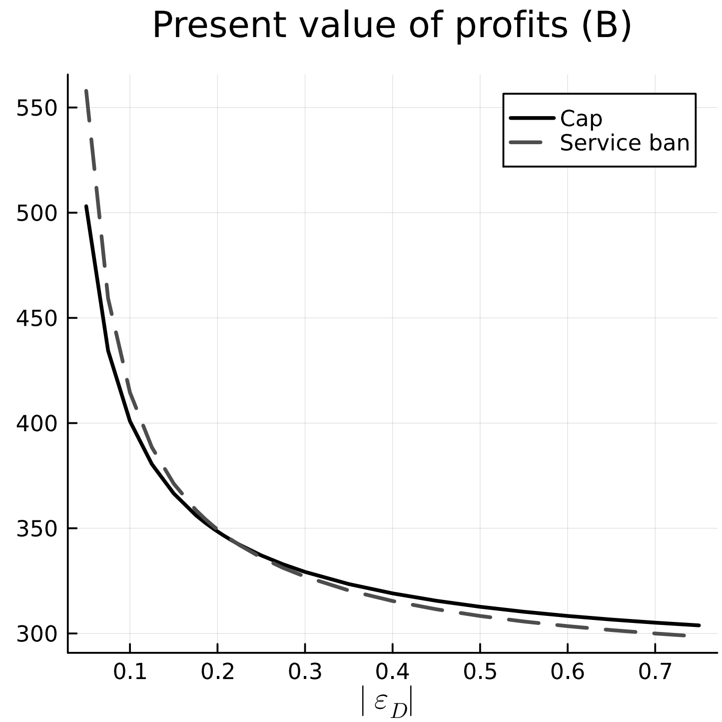

Figure A1: Present value of profits under different policy scenarios for a range of price elasticities of demand. Values are reported in billions of dollars in present value using the baseline discount rate (15% per year) until the sanction termination (80 quarters).

under the service ban or the cap, both policies continue to result in very similar profits in present value.

A.2 Discount Rate

Next, our analysis considers discount rates from 0 to 30% per year—twice as high as the baseline rate of 15% based on Russia's actual cost of capital. A very high discount rate, such as the one used for the upper bound of this analysis, can be interpreted as short-termism or myopia of Russia's decision making, implying that it would be placing an excessive weight on the first few quarters relative to the long run.

Figure A2 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. Here, we also recalibrate ϕ in each run to enforce the same calibration procedures used in the baseline. Negative values indicate that the PV of profits is lower under the baseline cap ($60) relative to a ban policy.[3]

[3] The simulations reported in Figure A2 use different discount rates for Russia. However, the displayed present

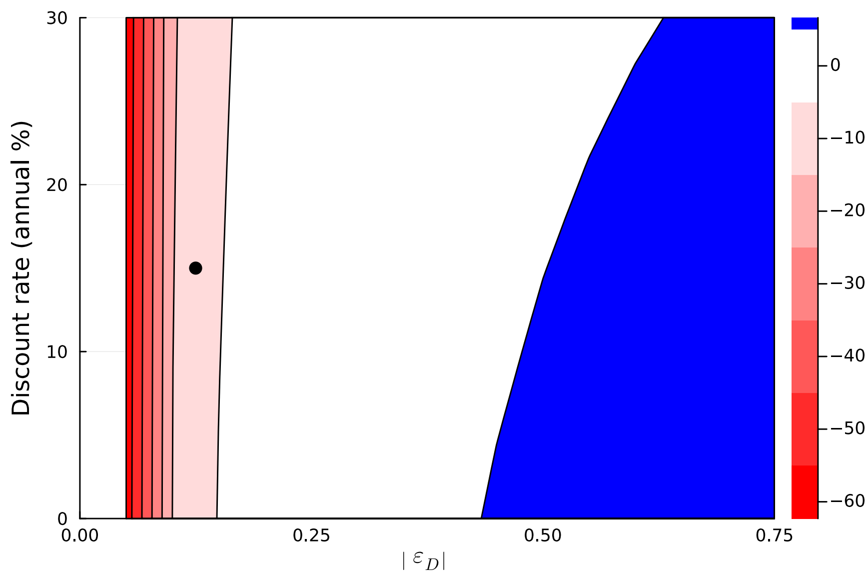

Figure A2: Difference in profits between the cap and the ban policies for various combinations of the discount rate and the price elasticity of demand. The black dot indicates the combination of parameters used in the baseline model. Values are reported in billions of dollars in present value discounted at the baseline discount rate (15% per year) until the sanction termination (80 quarters).

The almost vertical level curves in Figure A2 demonstrate that the ordering in the outcomes between the cap and ban policies is largely insensitive to the choice of the discount rate. Moreover, we highlight that for a wide range of combinations between discount rates and elasticities of demand, the absolute difference in the profits under both policies is below 5 billion dollars (the area in white). This result is primarily due to the calibration of the 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 proportionally 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 A3 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.

value of profits is calculated using a common rate to ensure fair numerical comparisons across different scenarios.

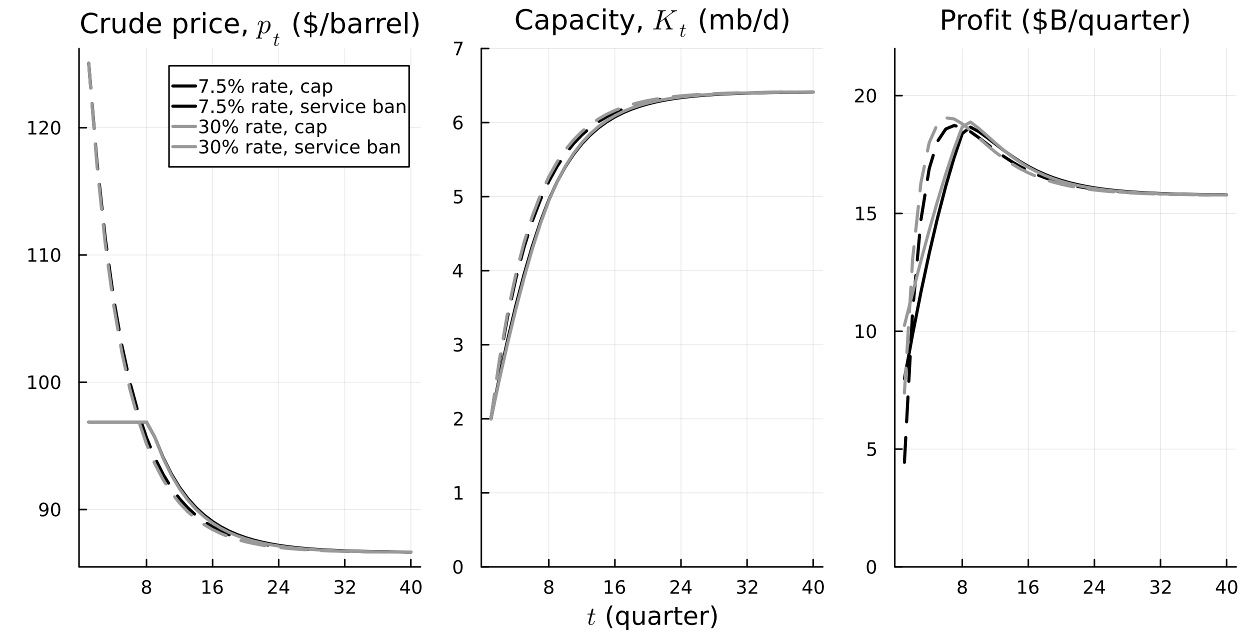

Figure A3: A comparison of the trajectories under a price cap sanction (solid) and a service ban (dashed). Black lines represent trajectories simulated using a 7.5% annual discount rate, whereas gray lines represent these trajectories using a 30% discount rate. All scenarios are based on the baseline price elasticity of demand ($\epsilon_D = -0.125$). These panels display only the first 40 quarters of the 80-quarter simulation.

A.3 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 A4, 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 A5 examines whether the relative ranking of profits under a service ban or cap flips with

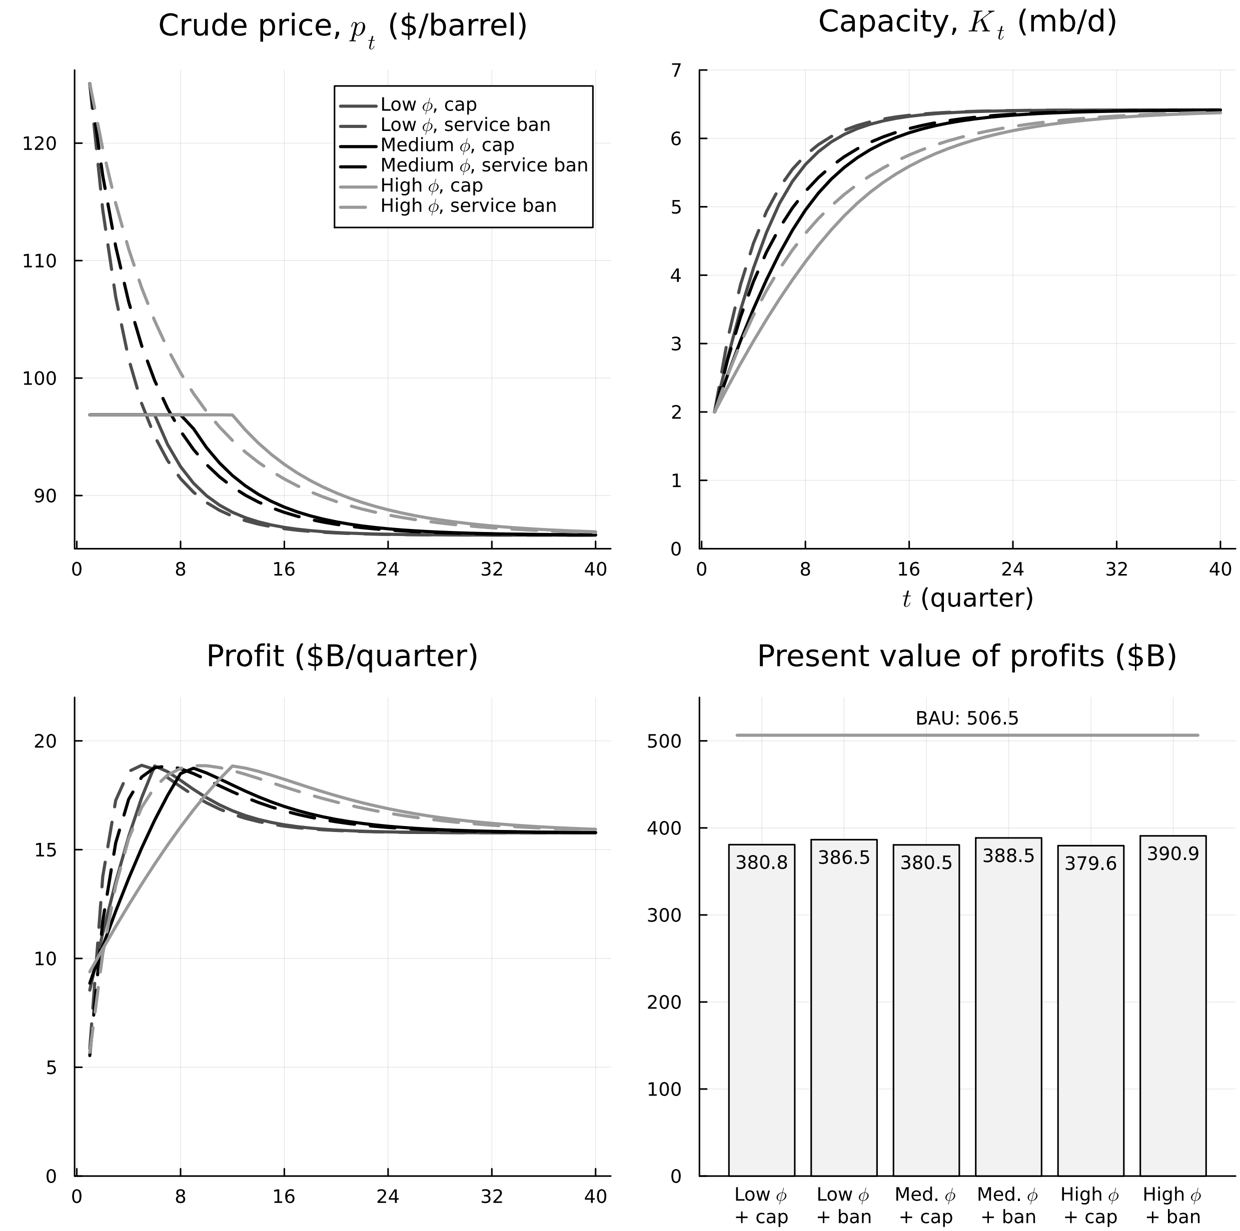

Figure A4: A comparison of prices, capacity, and profits under the service ban (dashed lines) vs. a \$60 price cap sanction (solid) with different marginal expansion costs (ϕ). The line graphs display only the first 40 quarters of the 80-quarter simulation.

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 ten 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.

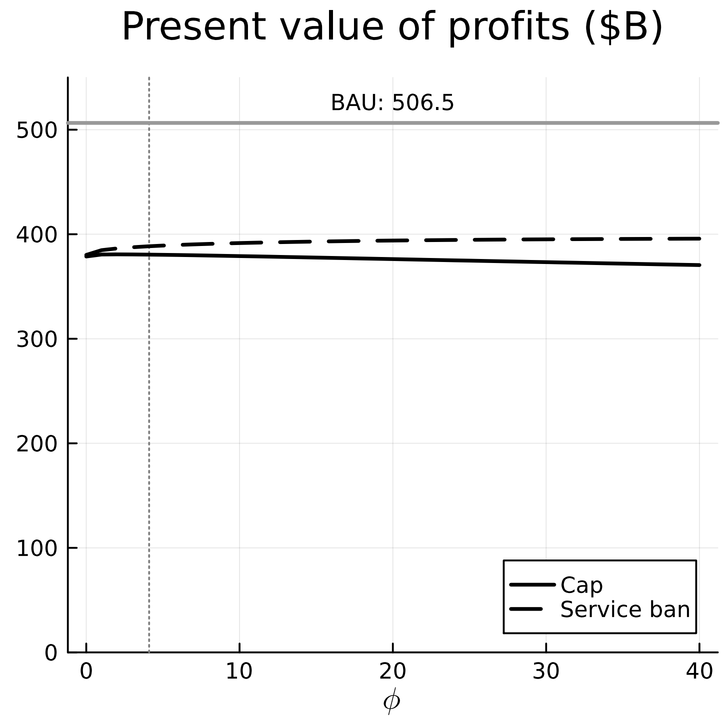

Figure A5: Present value of profits under the service ban or $60 price cap sanction for different levels of the marginal investment cost (ϕ parameter). The vertical dashed line in gray indicates the calibrated ϕ in the baseline model.

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 $C'(R^*) = P(Z_0 + R^*) - d$. 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 R^*. 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 represented in the diagram in Figure 1.

B Relaxing Baseline Assumptions

B.1 Endogenous Non-Russian Supply

In the baseline version of our model, we have assumed that non-Russian supply of oil is constant (at Z_0). If instead oil supply from the rest of the world was sensitive to the world oil price ($Z_0'(p) > 0$), we find that the price spike caused by the service ban is smaller and, as a result, the service ban might become more harmful to Russia than the cap although both sanctions deliver quantitatively similar profits in present value.

In this section, we examine how a non-zero price elasticity of non-Russian supply affects the outcomes of the cap and service ban. To do so, we maintain equations (1)–(4) but rewrite the market-clearing conditions as

$$p_t = P(Z_t + Q_t + X_t) \tag{B1}$$

$$Z_t = Z_0 \left(\frac{p_t}{p_0}\right)^{\epsilon_Z}, \tag{B2}$$

where $\epsilon_Z \geq 0$ is the assumed price elasticity and p_0 and Z_0 are defined in the baseline model (see Table 1). Thus, our baseline model is equivalent to setting $\epsilon_Z = 0$..

We simulate the extended model adding equations (B1–B2) for two positive values of ϵ_Z: 0.1 and 0.2. For reference, recent estimates of the price elasticity of global oil supply range from 0.01 (Kilian and Murphy, 2014) to 0.15 (Baumeister and Hamilton, 2019). As in the other extensions, we recalibrate ϕ (the slope of the marginal investment cost function) so as to generate capacity expansion trajectories comparable to observation and the baseline model.

Figure B1 displays the outcomes of these simulations. As shown in the top left panel, the more intense expansion of non-Russian supply with a higher ϵ_Z dampens the price spike under the service ban, as well as the level of the initial price plateau under the price cap. Moreover, with additional non-Russian supply, steady state prices are also slightly lowered relative to a fixed Z scenario. Following a lower price spike, incentives for capacity expansion are reduced, and the top right panel shows that a higher ϵ_Z induces both a lower steady state world price.

The impacts on Russian profits are displayed in the lower panels of Figure B1. The bottom left panel shows that as the price responsiveness of non-Russian supply increases, Russian profit

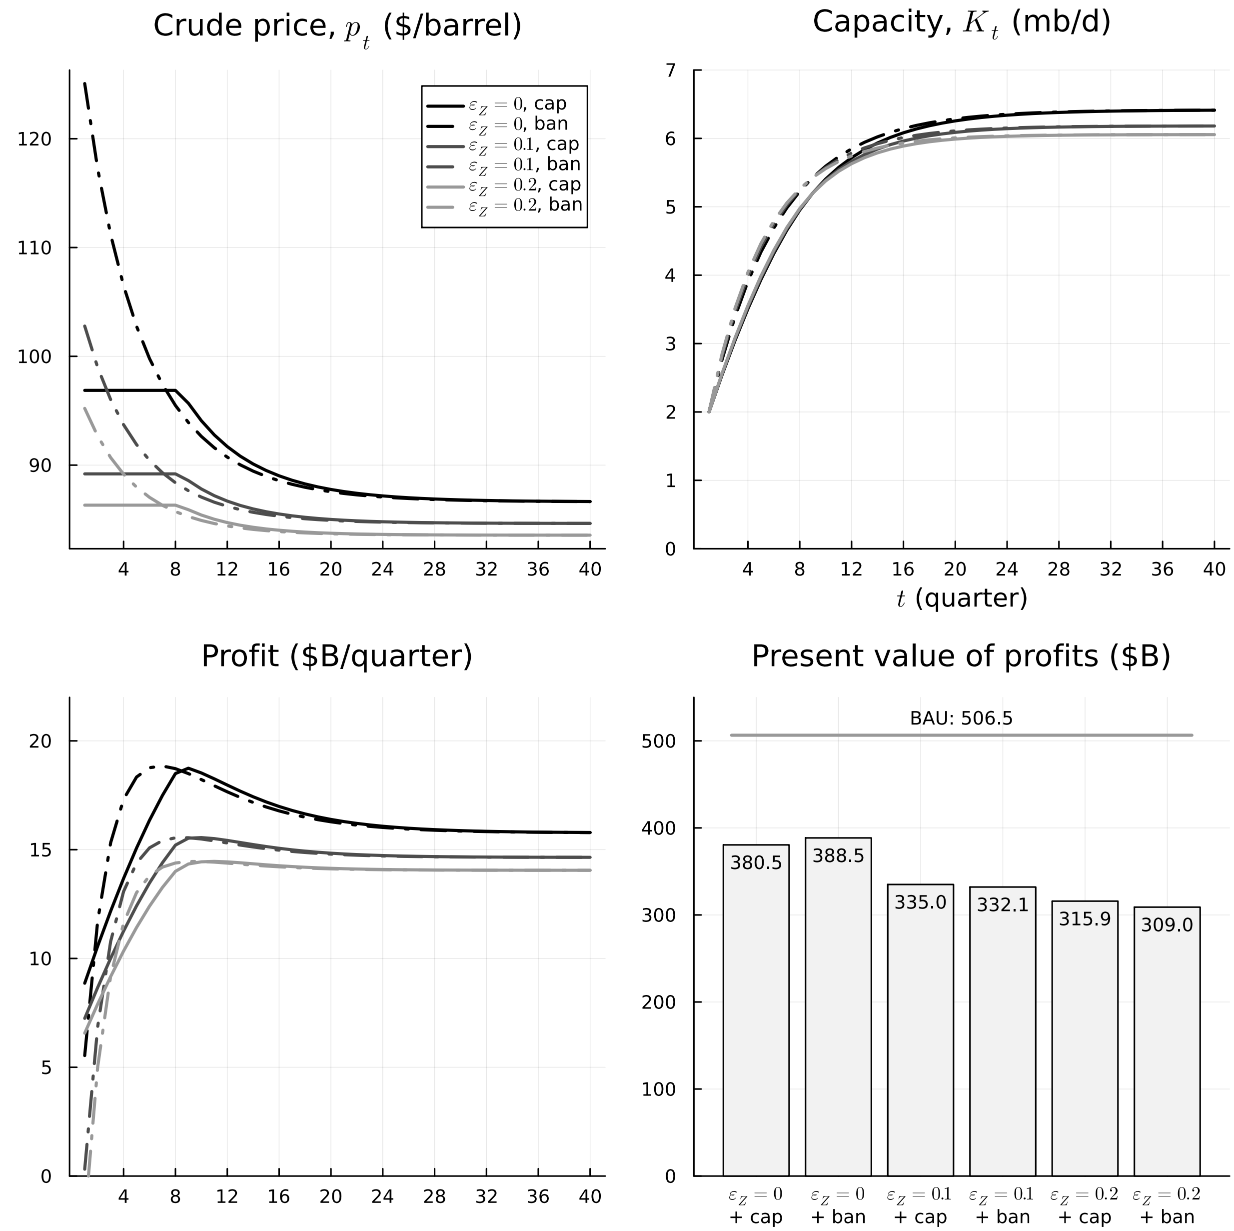

Figure B1: A comparison of the trajectories for different price elasticities of non-Russian export supply (ϵ_Z) when the sanction is a high price cap ($60) or a service ban. Each panel displays quarters 1–40 of the 80-quarter simulation.

opportunities shrink in the early periods, especially in the first eight quarters under the service ban. As a result, the ranking of PV of profits reverses for $\epsilon_Z = 0.1$ or 0.2. However, the present values under both policies remain close.

Next, we simulate both policy scenarios for intermediate values within the considered range of ϵ_Z. As Figure B2 reflects, the cap harms Russia's present value of profits more than the ban at

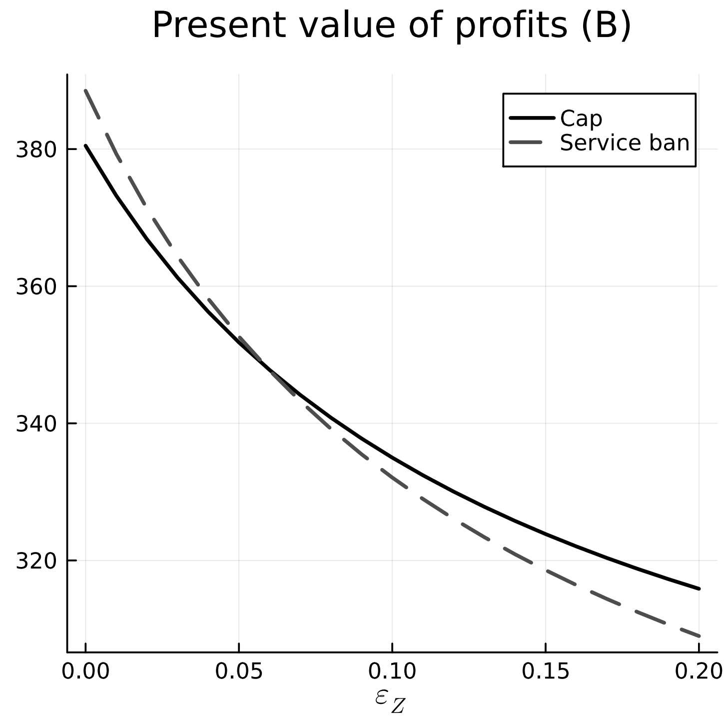

Figure B2: Present value of profits under the service ban and cap for a range of price elasticities of non-Russian supply ($\epsilon_Z \in [0.0, 0.2]$).

$\epsilon_Z = 0$. At the other extreme, the ban is more harmful. This graph also shows that the two policies are equally harmful at $\epsilon_Z \approx 0.06$.

B.2 Curvature of the Marginal Production Costs

Our baseline model assumes affine marginal costs. Nevertheless, empirical estimates of break-even prices of Russian oil assets indicate that the marginal cost curve may have significantly more curvature (Wachtmeister, Gars, and Spiro, 2023). In this Appendix, therefore, we generalize the marginal cost function to allow for varying degrees of curvature.

Previously, we calibrated the affine marginal cost function so it passed through two points. To allow for a more flexible functional form, we define

$$\tilde{C}'(X_t + Q_t) = c_0 + (p_0 - c_0) \left(\frac{X_t + Q_t}{Q_0} \right)^{\chi}, \tag{B3}$$

where c_0, p_0, and Q_0 are defined in the baseline model (see Table 1), and χ measures the curvature of this function passing through the same two points. $\chi = 1$ represents a linear function, whereas higher values of χ bend the curve towards a "hockey stick" form. Figure B3 depicts the marginal cost curve for various degrees of curvature. Note that the strictly convex marginal cost is below the linear marginal cost for aggregate exports below the initial level Q_0 (which is also the maximum in our simulations) of 7.4 mb/d.

We simulate the cap and the service ban policies with more convex marginal costs determined by values of χ at 2 and 8. Here, we also recalibrate the marginal investment cost function (parameter ϕ) in each of the cases to ensure that expansion trajectories under the cap are comparable with the baseline and across scenarios. The results are displayed in Figure B4. We note that a larger χ means that the marginal cost drops further for a given reduction in exports. Under competitive behavior, this means that supply is more inelastic to price changes. Therefore, Russian supply is above the baseline ($\chi = 1$) in all periods under the cap or after the first period under the service ban. The top panel in Figure B4 reflects these mechanisms. The top left shows that the world price drops more quickly under the ban and has longer plateaus under the cap, with both price paths converging to a lower steady state than in the baseline simulations. Similarly, the top right panel shows that a higher supply also means faster expansion under the ban that converges to a higher steady-state capacity, reflecting the lower marginal costs nearing Q_0.

The bottom panels of Figure B4 illustrate how profits change under different values of χ. Profits are generally higher, primarily because costs are lower. Since each scenario is simulated

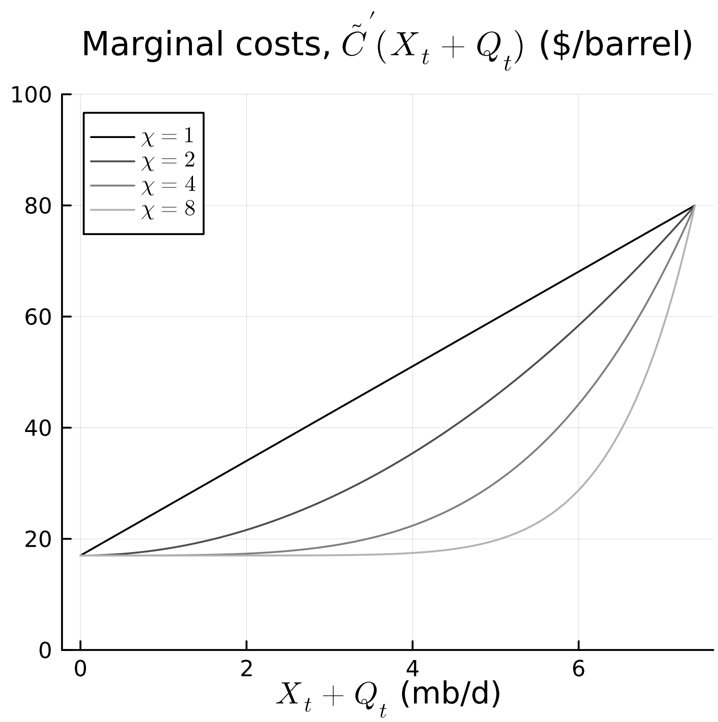

Figure B3: Varying curvature of marginal production costs with different levels of parameter χ.

under different costs, it is difficult to compare the present value of profits across scenarios. For this reason, in the bottom right panel we also plot what BAU profits would be under each χ. The vertical bars show that a higher convexity of marginal costs might lead to a higher present value of profits under the cap relative to the ban. This result is further illustrated in Figure B5, which shows that the ranking flips for $\chi > 2.36$. Nevertheless, both Figures B4 and B5 show that the present value of profits under each policy stays (i) close to each other and (ii) substantially lower than BAU, regardless of χ.

The main implication of χ is the fact that it modulates the supply responses to price changes. As Figure B3 shows, a higher χ leads to a steeper marginal cost curve near the initial export level Q_0. Although our stylized model is not equipped to capture and accurately predict the intricacies of global oil markets—much less so during the turbulent period following the Russian invasion—it is useful to at least gauge key generated outcomes with those observed in data. In particular, we examine average quarterly prices and supply as published in IEA's monthly Oil Market Reports (International Energy Agency, 2022, 2024b) in comparison with model outcomes.

Table B1 reports prices and quantities over different periods and the supply elasticity implied by their changes. In the observed data, we examine Urals prices and total Russian supply in the quarter immediately before the invasion (2021 Q4) and two years later (2023 Q4). This data shows that prices dropped by 9.6% between the end points of this comparison, whereas total Russian supply—including domestic and non-maritime exports—fell by about 5%. The ratio between these changes roughly implies an elasticity of 0.42. This magnitude is substantially above typical estimates of global oil supply (discussed in Appendix B.1), highlighting the exceptionalism of wartime dynamics and the restrictions imposed by the multiple sanctions on the Russian economy.

For the model results reported in Table B1, the initial period reflects data-driven assumptions: slightly rounding up Urals prices from 2021 Q4 and export volumes as reported in International Energy Agency (2024a). The final period is represented by the steady-state solution of the model in each case. Final prices in those cases represent producer prices, which are calculated as the world oil price minus the discount on shadow fleet sales. Here, we highlight how the cases of $\chi = 1$ (baseline) and $\chi = 2$ generate similar prices similar to those in the observed data, thus facilitating the comparisons. Since the model only simulates exports, we adjust implied elasticity calculations by multiplying back the export share of total supply volume. The rightmost column shows that implied elasticities for $\chi = 1$ indicate that this parameter choice generates stronger supply responses than those observed in the IEA data, whereas $\chi = 2$ best approximates those responses among the examined cases. Nevertheless, as indicated in our discussion above, qualitative results are not overturned by increasing $\chi = 2$, for which reason we maintain the parsimonious choice of a simpler functional form in the baseline.

	Initial price	Initial quantity	Final price	Final quantity	Implied elasticity
Observed					
IEA data	78.36	9.90	70.83	9.50	0.42
Model results					
$\chi = 1$	80.00	7.40	71.63	6.42	0.95
$\chi = 2$	80.00	7.40	69.34	6.75	0.50
$\chi = 4$	80.00	7.40	67.59	7.01	0.26
$\chi = 8$	80.00	7.40	66.44	7.18	0.13

Table B1: A comparison of implied elasticities. *Observed* data reports quarterly average Urals prices in 2021 Q4 (initial) and 2023 Q4 (final) in dollars per barrel, and total supply quantities in mb/d for the same periods. *Model results* rows present outcomes under different χ values. For model results, the initial period is $t = 0$, before any sanctions, and the final period is the steady state solution at $t = T$; price columns report producer price (world price minus the shadow fleet discount) and quantities represent exports in mb/d. The implied elasticity is the ratio of percent changes in quantities and prices, adjusted for the export share of total supply ($\approx 74\%$) in the model results.

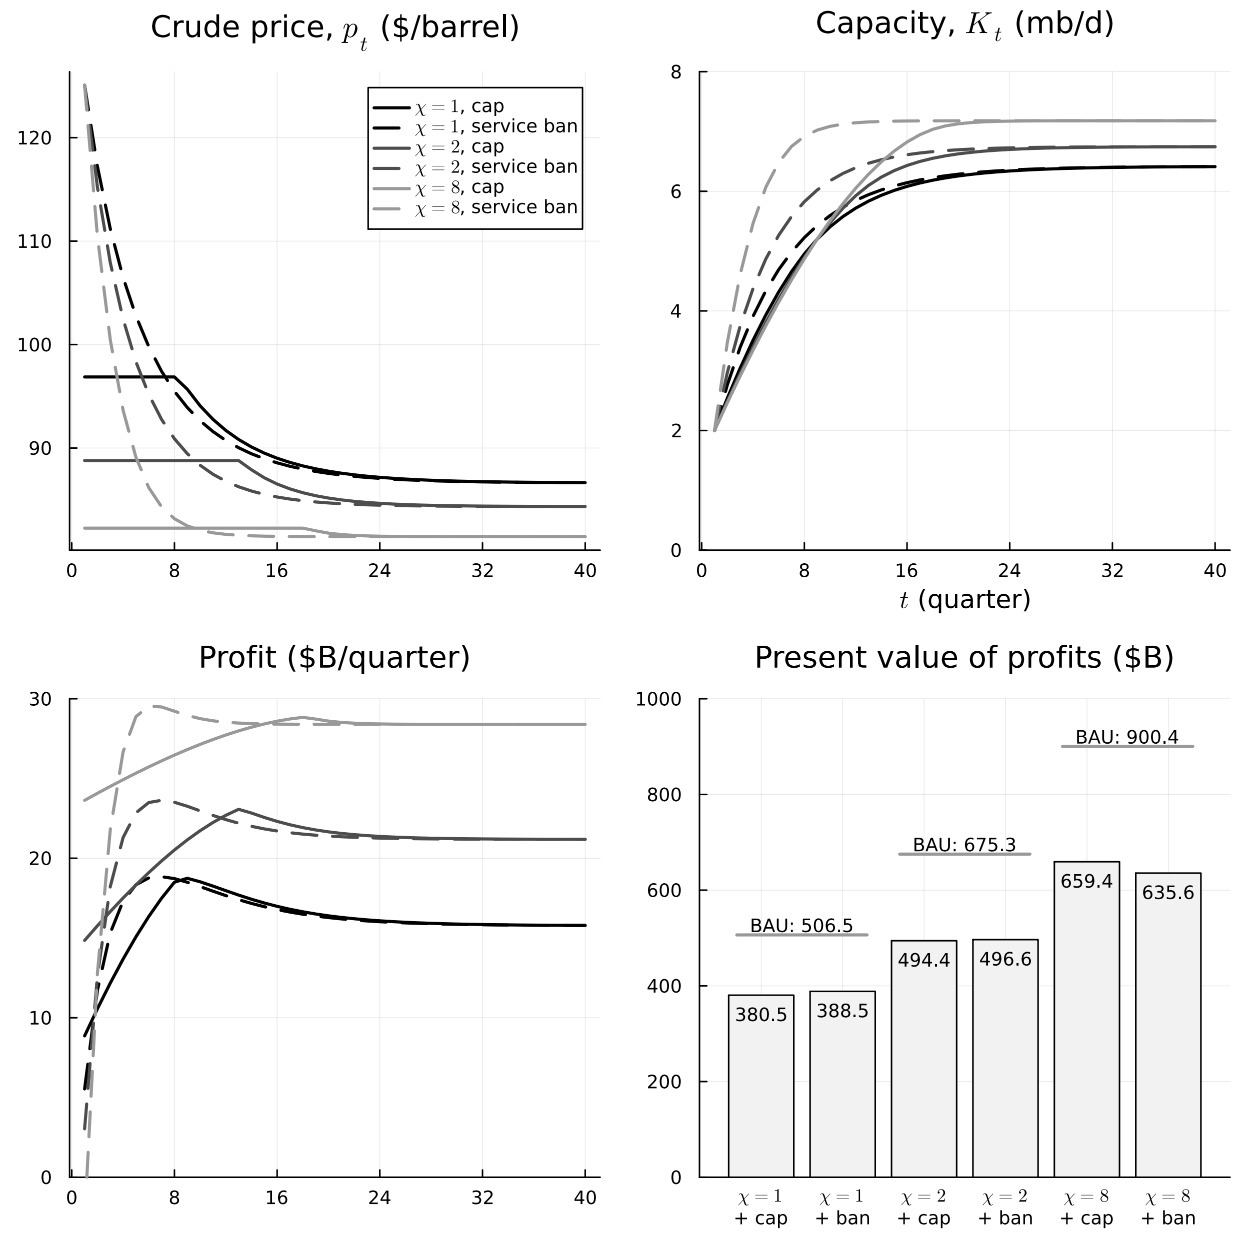

Figure B4: A comparison of the trajectories for different curvatures of the marginal production cost function under a price cap or a service ban policy. $\chi = 1$ defines a linear marginal cost curve, and higher value increase its convexity. Each panel displays quarters 1–40 of the 80-quarter simulation.

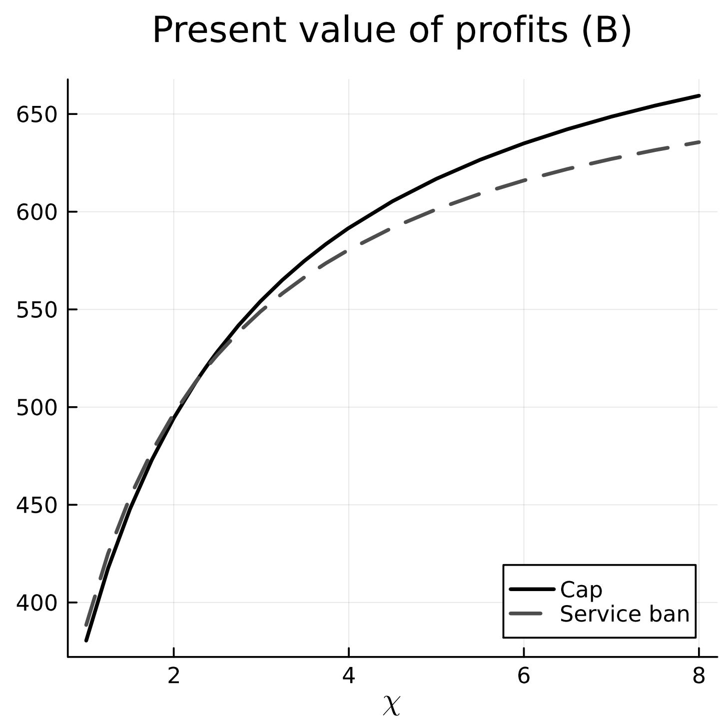

Figure B5: Present value of profits under the service ban and cap for a range of marginal cost curvatures ($\chi \in [1, 8]$).

B.3 Endogenous Discount on Shadow Fleet Sales

In our baseline model, the discount at which Russia sells its exports via the shadow fleet is held fixed at $d = 15$. As indicated in the main text, this value approximates the differential between Brent and Urals reference prices when it stabilized about one year after the start of sanctions.

Besides an insurance premium and higher shipping costs, the literature suggests this discount could reflect monopsony power and a risk of being found in violation of the cap. Arguably, the opportunity cost of Russia having to sell at the capped price if negotiations with importing countries broke down might account for the bargaining power of non-coalition importers. With imperfect enforcement, if the audit probability were increased, the expected price per barrel would fall, less cheating would occur, and the world price would rise, so that the discount on Russian exports would rise due to the risk of being found in violation of the cap.

This subsection extends our baseline model to include an endogenous discount on shadow fleet sales. In this extension, the maximum value of evasion is defined by the arbitrage opportunity, which is the difference between the world oil price and either the price cap (for cap scenarios) or the marginal production cost (for ban scenarios). We posit that d represents the portion of the arbitrage value that is absorbed by agents other than Russia (monopsony power from new buyers, transportation, increased insurance costs and risk premia, etc). Based on this, we parametrize an endogenous discount as a function of the arbitrage opportunity offered by evasion,

$$d\left(P(Z_0 + K_t) - \hat{p}, K_t\right) = \nu_0 + \nu\left[P(Z_0 + K_t) - \hat{p}\right], \tag{B4}$$

where ν_0 and ν are the parameters to be determined. We highlight that, in equilibrium, K_t itself is a function of $\hat{p}$.

Following our choice of $d = 15$ based on data, we interpret this value as the steady state discount, which is obtained when $Q_t = 0$ and $X_t = R^*$. We note that Kilian, Rapson, and Schipper (2024) estimated added transportation costs using the shadow fleet were between \$12 and \$15 per barrel. Therefore, our interpretation could be understood as transportation costs remaining as the long-lasting component of the discount over time. Moreover, we observe that in the weeks following the invasion (March and April of 2022), the gap between Urals and Brent prices achieved its peak, at approximately \$36 per barrel; during the same period, the Brent price peaked at around \$123. In our model, the value of evasion is the biggest in the initial period of the service ban scenario, in which world price peaks and shadow fleet exports are limited to K_1. We use these two observations—the peak shortly after the invasion and the steady state level—to establish the following conditions

$$\begin{aligned} \nu_0 + \nu\left[P(Z_0 + R^*) - \hat{p}\right] &= 15 \\ \nu_0 + \nu\left[P(Z_0 + K_1) - C'(K_1)\right] &= 36. \end{aligned}$$

For reference, in our calibrated baseline model we have $R^* = 6.4$ mb/d, $P(Z_0 + R^*) = \$86.62$, $P(Z_0 + K_1) = \$125.07$, and $C'(K_1) = \$34.03$. Solving for the parameters, we obtain $\nu_0 = 6.32$ and $\nu = 0.33$.

We build on this idea to expand the model and allow for an endogenous discount that varies with the arbitrage opportunity offered by shadow fleet sales. In doing so, we revise equation (2). Since $p_t - d_t = p_t - [\nu_0 + \nu(p_t - p_0)]$,

$$X_t \geq 0, [(1-\nu)p_t + \nu\hat{p}] - \nu_0 - C'(Q_t + X_t) - \alpha_t \leq 0, \ c.s.. \tag{B5}$$

We simulate the effects of a \$60 cap and a service ban using this extended model. Figure B6 shows the results of these simulations. As in most of our extensions, we recalibrate the marginal investment cost function for the endogenous discount case to generate capacity trajectories under the cap that are comparable with the baseline and across scenarios. Due to this recalibration, the top panel shows that capacity and price trajectories under the cap perfectly overlay, regardless of the discount model (solid lines). Also for the cap policy, the bottom left panel shows that the endogenous discount model leads to a slightly higher discount during the "price plateau" phase, but ultimately converges to the fixed discount level. This difference results in a small decrease in the present value of profits relative to the baseline model with a fixed discount, as shown in the bottom right panel.

Substantial differences in trajectories occur, however, in the service ban scenario. With the absence of a channel to export at the cap, the value of the arbitrage opportunity is amplified. As the bottom left panel in Figure B6 shows, the endogenous discount in the service ban case sits at a higher level than the fixed baseline value of \$15 throughout the entire simulation. Notably, it converges to approximately \$25 per barrel. This wedge lowers the incentive to capacity expansion, resulting in a lower steady-state fleet (top right panel) and a higher steady-state world price (top left panel). Consequently, with fewer exports and lower effective export prices, Russian profits under a service ban with endogenous discount are more negatively impacted relative to the cap scenario. The drop in the present value of profits with endogenous discount is enough to flip the ranking of profits and create a significant gap between both outcomes—and both remain substantially below BAU profits. Further simulations holding $\nu_0 = 6.32$ constant and varying ν show that both policies result in equal harm to Russian profits with a value of $\nu \approx 0.018$.

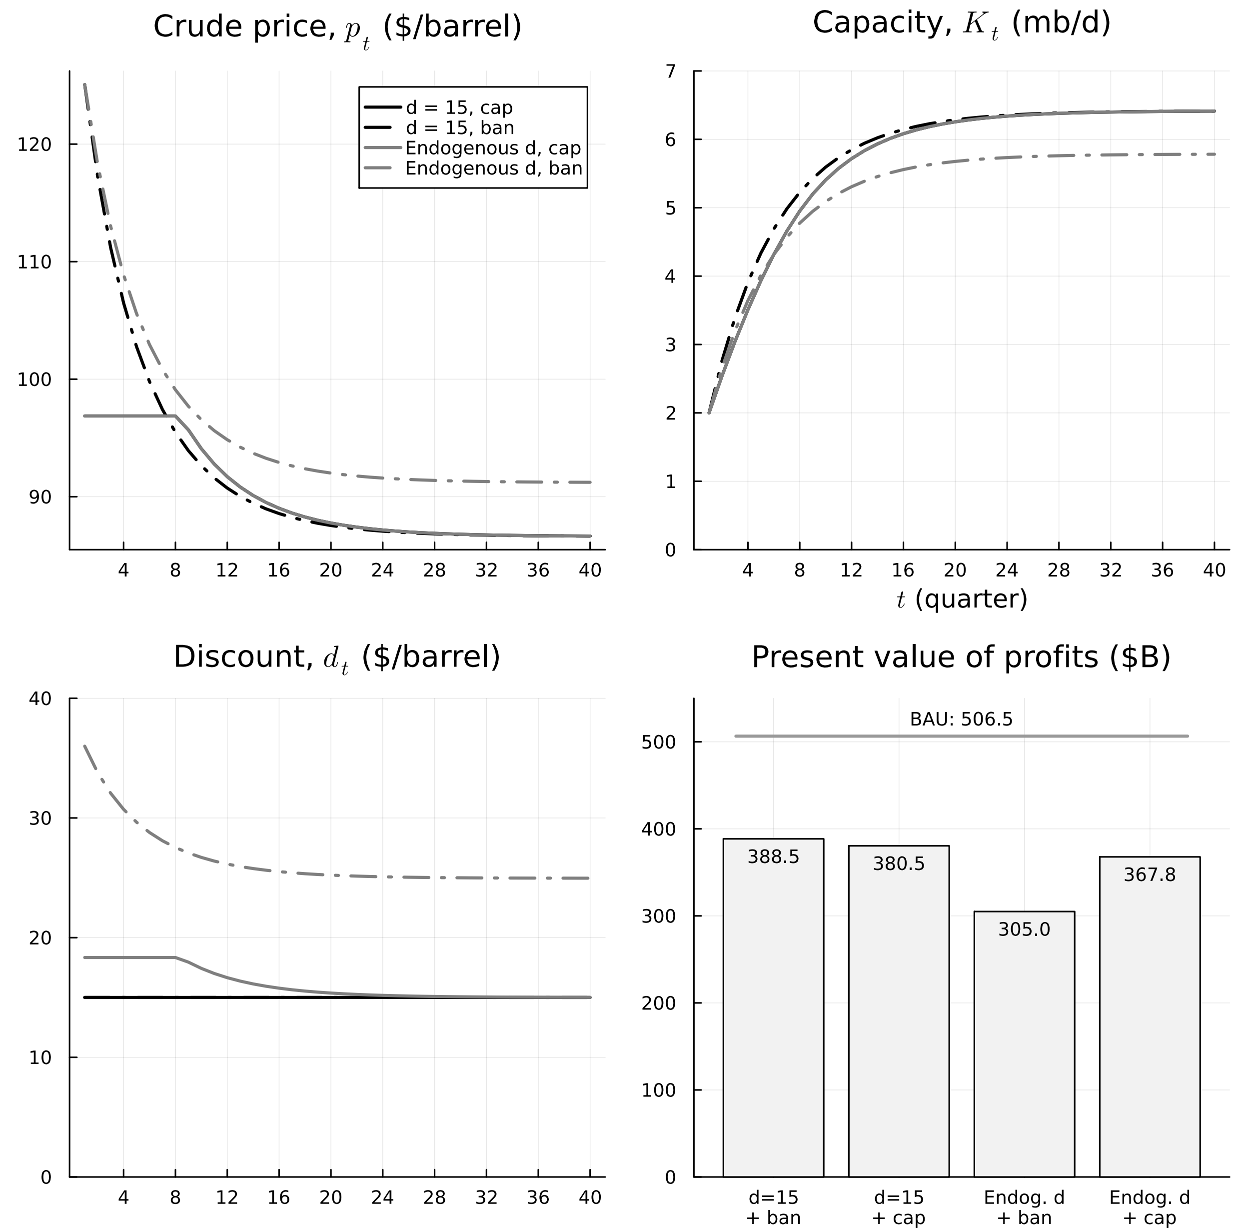

Figure B6: A comparison of the trajectories under a price cap or a service ban policy based on fixed vs. endogenous shadow fleet discount models. Each panel displays quarters 1–40 of the 80-quarter simulation. On the top panels, the curves for “d = 15” and “Endogenous d” perfectly overlap due to the recalibration of the investment cost function.

B.4 Internalization of Price Effects

A paradoxical result in our dynamic model is that Russia, a profit-maximizer, can benefit from a lower price cap even though Russia could have seemingly secured the same benefits without any change in the cap. This result is not an artifact of our price-taking assumption. For example, Turner and Sappington (2024) reach a similar conclusion in their static Cournot duopoly, although, as discussed in our literature review, our results are very different from theirs. Yet, as mentioned in the introduction, such a paradox would not arise if Russia were the only agent fully internalizing its market influence, behaving as a monopolist facing the demand for its oil. In reality, as explained in Section 6, Russia's oil supply behavior has been more consistent with the price-taking hypothesis than with the exercise of monopoly power, partly because of the structural imperfect coordination of Russian oil suppliers. In this appendix, we extend the baseline model to examine cases in which Russia internalizes the intertemporal price effects at various degrees.

We have two objectives: (1) to quantify the present value of Russian profits as a function of the cap for different levels of price effect internalization; and (2) to show the invariance of the present value of Russian profits as the level of internalization approaches the monopoly case.

We modify the problem in Subsection 2.3 by replacing $\{p_t\}_{t=1}^T$ with $P(Q_t + X_t + Z_0)$. As a consequence, the first-order conditions (1) and (2) are modified as follows

$$Q_t \geq 0, \qquad \hat{p} + \theta X_t P'(Q_t + X_t + Z_0) - C'(Q_t + X_t) \leq 0, \ c.s. \tag{B6}$$

$$X_t \geq 0, \qquad P(Q_t + X_t + Z_0) - d + \theta X_t P'(Q_t + X_t + Z_0) - C'(Q_t + X_t) - \alpha_t \leq 0, \ c.s. \tag{B7}$$

The remaining equations, (3) and (4), are unchanged. The term $X_t P'(Q_t+X_t+Z_0)$ in conditions (B6-B7) 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. However, in both conditions, parameter $\theta \in [0, 1]$ modulates the degree of internalization of price effects. When $\theta = 0$, the agent behaves as a price taker and does not take into account price responses to changes in supply; when $\theta = 1$, the agent fully internalizes its monopolist position. Values between 0 and 1 represent partial internalization of price effects, which could result, for example, from the lack of internal coordination among Russian producers. In practice, this intermediate case will most likely approximate Russia's aggregate supply decisions, as the hybrid structure of the oil production sector in Russia leads to decentralized decisions and players competing for exports (Henderson and Fattouh, 2016).

We simulate the solutions for $\theta \in \{0, 0.25, 0.5, 0.75, 1\}$ and cap levels between \$30 and \$70 per barrel.[4] This exercise assumes that any price internalization behavior starts at $t = 1$. All simulations use our baseline calibration of the investment cost function F. Unlike other sensitivity analyses and extensions, we do not recalibrate parameter ϕ (the slope of F) because for $\theta \geq 0.5$, there are no values of ϕ that can make the model match observed capacity expansion.[5]

Figure B7 presents the value of profits Russia earns under each internalization degree at different price caps. As expected, the present value of Russian profits is higher under monopoly ($\theta = 1$) than under price taking ($\theta = 0$). For a high degree of internalization ($\theta \geq 0.75$), the present value of profits does not increase as the cap decreases. However, with lower degrees, there is a slight decrease in profits as the cap increases, with the most pronounced cases occurring when θ approaches 0. Despite these differences, we note that a key qualitative finding still stands: The present value of profits under a low cap (30) and a high cap (70) are very similar, regardless of the degree of internalization.

Despite similar results for the present value of profits under either policy, this exercise indicates that changes in θ can affect their ranking. We further investigate this notion in Figure B8, which displays the present value of profits for either policy for $\theta \in [0, 1]$. This plot can be seen as representing how profits vary on a continuum of θ along two vertical slices in Figure B7: one at \$30 and another at \$60. Again, this plot illustrates how close the present values of profits are under each policy, regardless of the level of price effect internalization. Numerical inspection of these simulations shows that both lines cross at approximately $\theta = 0.85$, indicating that profits under a \$60 cap would slightly exceed those under a service ban only under high degrees of internalization of price effects.

To further demonstrate how the different degrees of internalization impact profits, Figure B9 displays the trajectories of prices, capacity, and marginal markups under three scenarios: no internalization ($\theta = 0$), partial internalization ($\theta = 0.5$), and full internalization ($\theta = 1$). Based on

[4] For 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.

[5] This is because higher degrees of internalization always lead to a steady state capacity below the observed level.

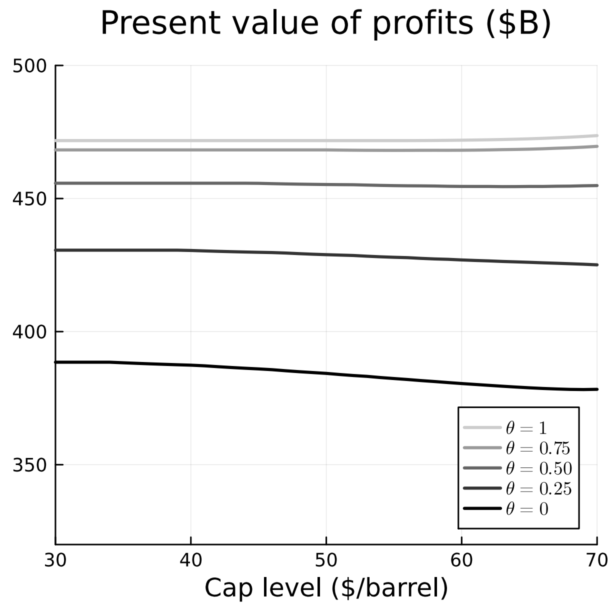

Figure B7: A comparison of profits for different degrees of price effect internalization (θ) under various levels of the price cap. Here, $\theta = 0$ indicates no internalization of price effects, whereas $\theta = 1$ indicates full internalization.

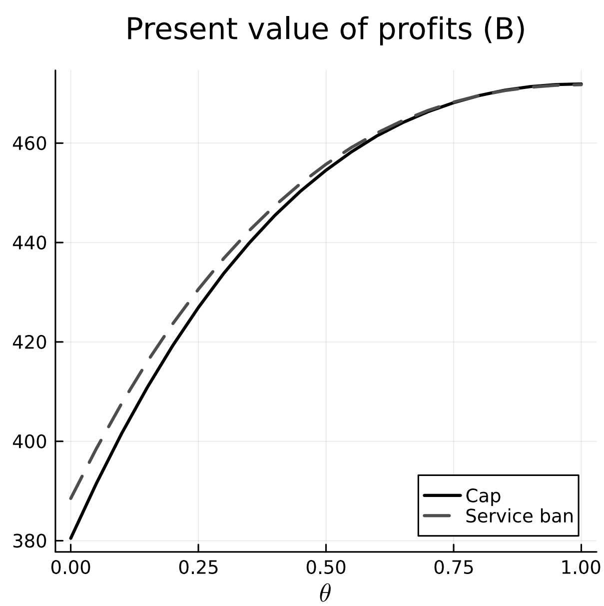

Figure B8: Present value of profits under different policy scenarios for a range of degrees of internalization of price effects. Here, $\theta = 0$ indicates no internalization of price effects, whereas $\theta = 1$ indicates full internalization.

equation B7, marginal markups are defined as the gap between marginal revenue and marginal costs (accounting for the shadow value of a constrained fleet):

$$\mu_t \equiv P(Q_t + X_t + Z_0) - d - C'(Q_t + X_t) - \alpha_t = \theta X_t P'(Q_t + X_t + Z_0). \tag{B8}$$

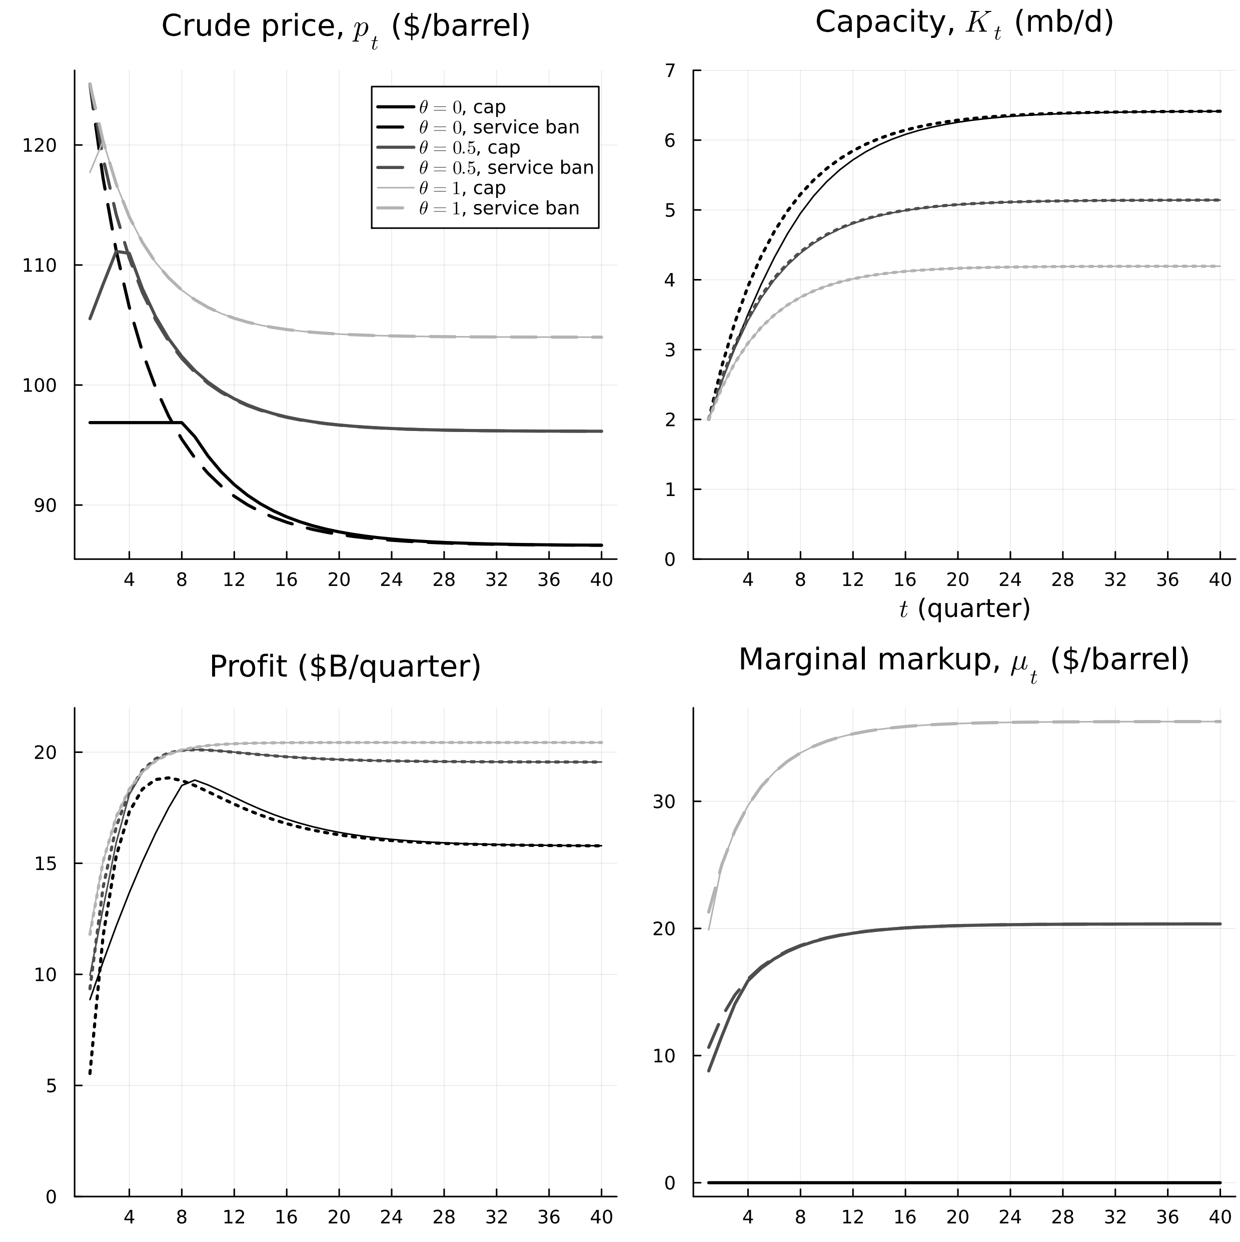

Figure B9: A comparison of the trajectories for different degrees of price effect internalization (θ) when the sanction is a high price cap ($60) or a low price cap (equivalent to a service ban). For $\theta = 1$, the two trajectories in each panel are indistinguishable since they *coincide*. Each panel displays quarters 1–40 of the 80-quarter simulation.

The panels in Figure B9 illustrate the general intuition that a greater internalization of price effects leads to lower supply—in the form of more limited fleet expansion—and sustained higher prices and markups. We note that full internalization trajectories in each of the four panels for the complete service ban and the $60 cap are *identical* and coincide. The only exception is the price (upper left panel) in the very first quarter. Being unable to enlarge 1^{st}-quarter shadow fleet capacity, Russia would sell a small portion of exports using Western services ($Q_1 \approx 0.7$ mb/d). However, after the 1^{st} quarter, it never uses Western services again, and the price trajectory thereafter coincides with its response to a service ban. For both policies, if Russia fully internalizes its price effects, it expands the shadow fleet far less than a price-taker would. As a result, the world price and the profits are uniformly higher.

While the outcomes for the case of partial internalization generally stand between the extreme cases, the quarterly profits trajectory (bottom left panel) illustrates an important mechanism. While full internalization leads to a monotonic increase in profits, a less-than-perfect internalization causes profits to peak. Further inspection shows that profits peak precisely when capacity reaches the steady-state level of $\theta = 1$. However, due to an underestimation of price effects under partial internalization, the agent continues to expand capacity, thus lowering quarterly profit levels.

B.5 Delayed Implementation and Anticipatory Behavior

Our baseline model assumes the immediate implementation of sanctions, with the shadow fleet expansion occurring while sanctions are in effect. The baseline parametrization of an initial evasion capacity $\bar{k} = 2$ mb/d is based on pre-invasion levels and reflects the impacts of policies implemented before Russia expanded its shadow fleet. However, in reality, the \$60 cap on Russian crude oil came into effect on December 5, 2022—a little over nine months after the invasion of Ukraine. By the time this sanction was implemented, Russia was already amassing its shadow fleet and diverting exports, most notably to India, China, and Turkey. At the end of 2022, Russian average exports to the "non-coalition" regions (China, India, Turkey, and the Middle East) had increased by 1.4 mb/d (International Energy Agency, 2024a). This anticipatory behavior gave Russia a head start, investing in evasion capacity and lowering the effectiveness of policies the West implemented later.

In this subsection, we extend the baseline model to examine the impact of delaying the implementation of sanctions on Russian profits. To do so, we assume that no sanctions are implemented in the first three quarters. Then, a service ban or a \$60 cap becomes effective in the 4^{th} quarter. We assume Russia anticipates this schedule and, taking the equilibrium price path as given, optimizes. Figure B10 displays the results of these simulations and compares the outcomes of immediate vs. delayed implementations.

We focus first on the comparison between delayed and immediate implementation of the cap (solid lines). The top left panel in Figure B10 shows that, with a delayed implementation, the world price stays at $p_0 = 80$ for three quarters. The price then jumps up to the plateau and stays at that level until the 8^{th} quarter. Oil prices begin to decline in the 9^{th} quarter for both the delayed and immediate cap implementation scenarios, indicating that cap sales end at that point, regardless of whether there is a three-quarter delay in implementation or not. The decline in prices, however, is slightly slower under a delayed implementation of the cap. This leads to tiny differences in price trajectories, with the biggest price gap of \$0.54 happening in the ninth period. These differences are due to a slightly slower capacity expansion when the implementation is delayed.

To further investigate these small differences in capacity additions, we report the values for the first six quarters in Table B2. Under a delayed implementation of the cap, the returns of an early expansion accrue later. Time discounting then leads to lower investment in the first quarters, with investment growing at the rate of $1/\beta$ ($\approx$4.15%) per quarter in the first three quarters. This mechanism leads to small differences in capacity trajectories, as shown in the top right panel of Figure B10, where the solid curves almost overlay. The cumulative investment in the first three periods is smaller under a delayed cap (1.38 mb/d added) compared to an immediate cap (1.51 mb/d). We note that this simulated expansion is closely aligned with the expansion of 1.4 mb/d in Russian exports to "non-coalition" regions observed in the IEA data.

As a result of this anticipatory expansion, when the delayed cap comes into effect in the fourth quarter, Russia has a shadow fleet capacity of $K_4 = 3.38$ mb/d. This increased evasion capacity softens the more substantial harm of sanctions in the early period, as shown in the bottom left panel of Figure B10. Under delayed implementations, Russia has a higher producer surplus in the first three quarters, and the downward trajectory of profits is the result of increasing investment costs in preparation for the sanctions. When the delayed cap becomes effective, the trajectories of profits under each policy remain close to their immediate implementation counterparts. The bottom right panel in Figure B10 shows that the ability to expand capacity in anticipation of sanctions raises the present value of Russian profits under the delayed cap relative to its immediate implementation. However, despite the significant head start to adjust, these gains are modest: about 5% (service ban) higher profits in present value, with a still substantial gap to BAU profits.

Comparing the immediate and delayed implementation of a service ban, we observe that small differences in trajectories (dashed lines) are due to the same incentives that lead to slightly different investments. The top left panel in Figure B10, in a delayed service ban implementation, the world price jumps up in the 4^{th} quarter, staying slightly above its immediate implementation counterpart as they decline until convergence around the 20^{th} quarter. Once again, these small differences reflect a slower capacity expansion under a delayed service ban, as shown in the top right panel. Table B2 complements this graph by showing that, when the ban comes into effect, Russia would have amassed enough capacity to export 3.65 mb/d. This anticipatory expansion avoids the profit loss in the first three quarters that would have occurred with an immediate implementation, as shown in the bottom left panel in Figure B10. As a result, however, the present value of profits under a delayed service ban is only 3.8% higher than under an immediate sanction.

Finally, these simulations assume Russia had perfect foresight of sanctions, whereas, in reality, the form of sanctions to be ultimately implemented was uncertain. Incorporating expectations over uncertain policy stringency would require a substantially more complex model. However, such an addition would be unlikely to change these results significantly because we observe that expansion trajectories are very similar, whether Russia expects the most stringent policy under consideration (service ban) or a less stringent policy (cap).

	Scenario			
Variable	Immediate cap	Delayed cap	Immediate service ban	Delayed service ban
I_1	0.54	0.44	0.76	0.53
I_2	0.50	0.46	0.63	0.55
I_3	0.47	0.48	0.52	0.57
I_4	0.43	0.44	0.43	0.47
I_5	0.38	0.40	0.35	0.39
I_6	0.34	0.35	0.29	0.32
$I_1 + I_2 + I_3$	1.51	1.38	1.91	1.65
K_4	3.51	3.38	3.91	3.65

Table B2: A comparison of the trajectories of investments (capacity additions), in the first six quarters. The two bottom rows show the cumulative investments until the 3^{rd} quarter and the capacity starting in the 4^{th} quarter, given the initial capacity of 2 mb/d. All values are mb/d.

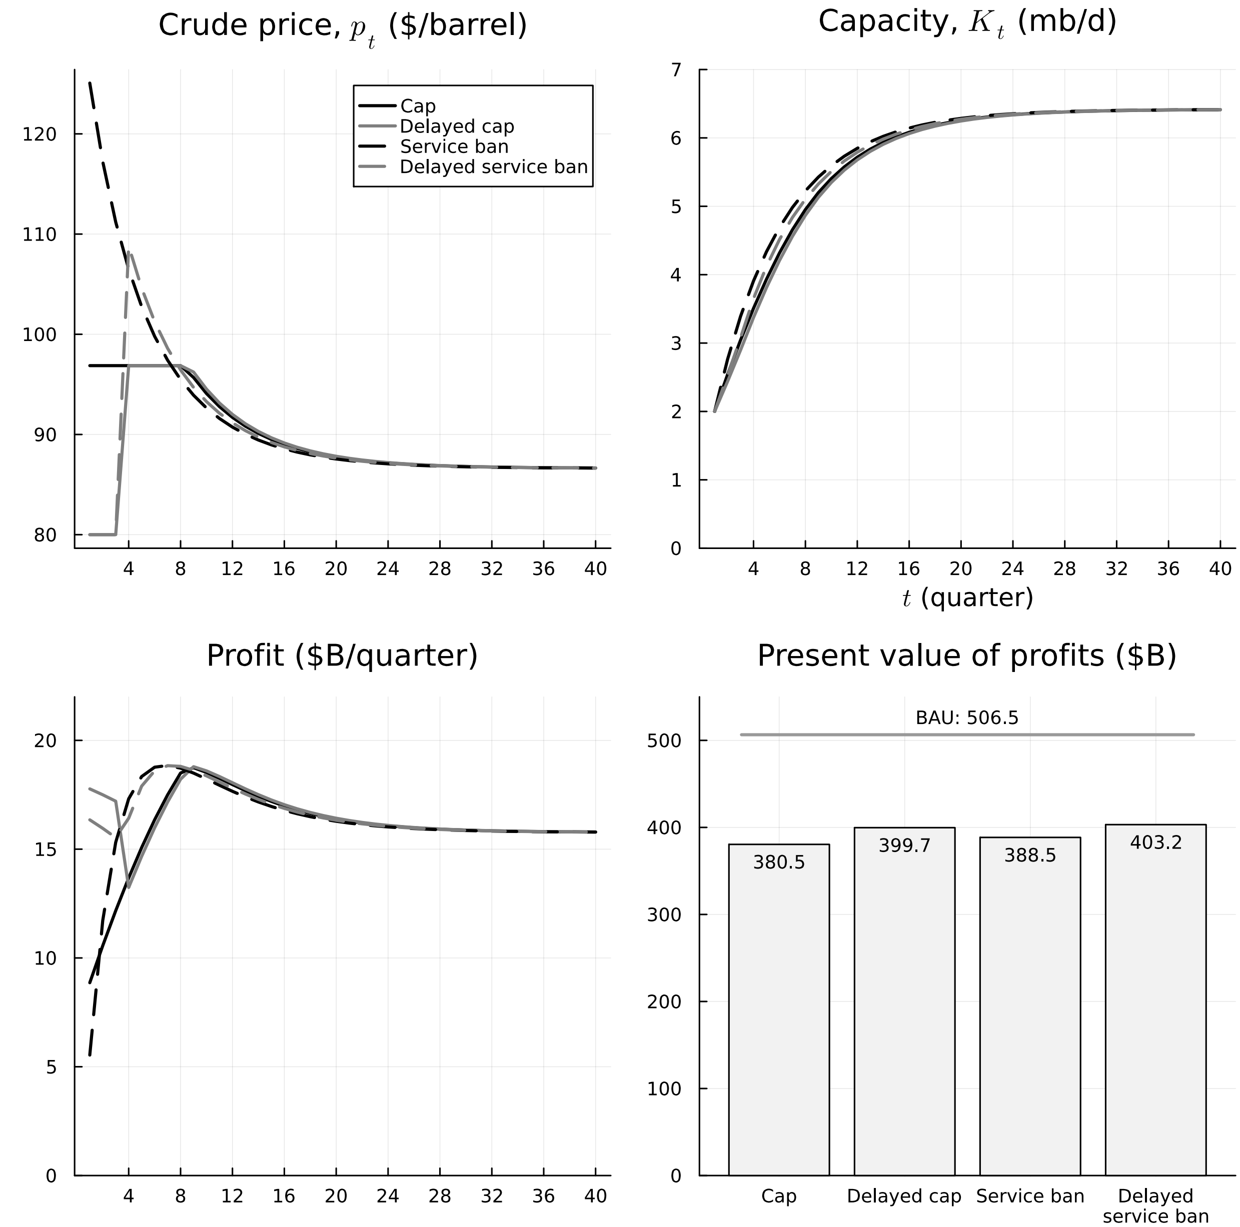

Figure B10: A comparison of the trajectories under a price cap or a service ban policy implemented either immediately or delayed by four quarters. Each panel displays quarters 1–40 of the 80-quarter simulation.

C Targeting the Shadow Fleet while Changing the Enforcement of the Price Cap

The following simulations extend those presented in Subsection 5.3 of the main text, entitled "Unanticipated Targeting of the Shadow Fleet" to the case of sanctions targeting the shadow fleet paired with changes in the enforcement policy.

We maintain the \$60 cap and simulate unanticipated targeting, examining the effect of an unexpected loss in the shadow fleet in the 12^{th} quarter that reduces the size of the shadow fleet to the level Russia had attained 8 quarters earlier: from K_{12} to K_4.[6]

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 12^{th}-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.[7]

Consider first the case where enforcement is initially lax. The solid lines in Figure C11 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 post-targeting present value of Russian expected profits (i.e., discounted to $t = 12$). In the bottom right panel of Figure C11, the reduction in the post-targeting 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 36^{th} quarter instead of the 28^{th} 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 simultane-

[6]Note that Subsection 5.3 considers targeting at K_{12} that reduces capacity to the level attained 4 quarters earlier. Therefore, the targeting reported in this Appendix represents a larger capacity reduction.

[7]With $a = 0.2$ 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%.

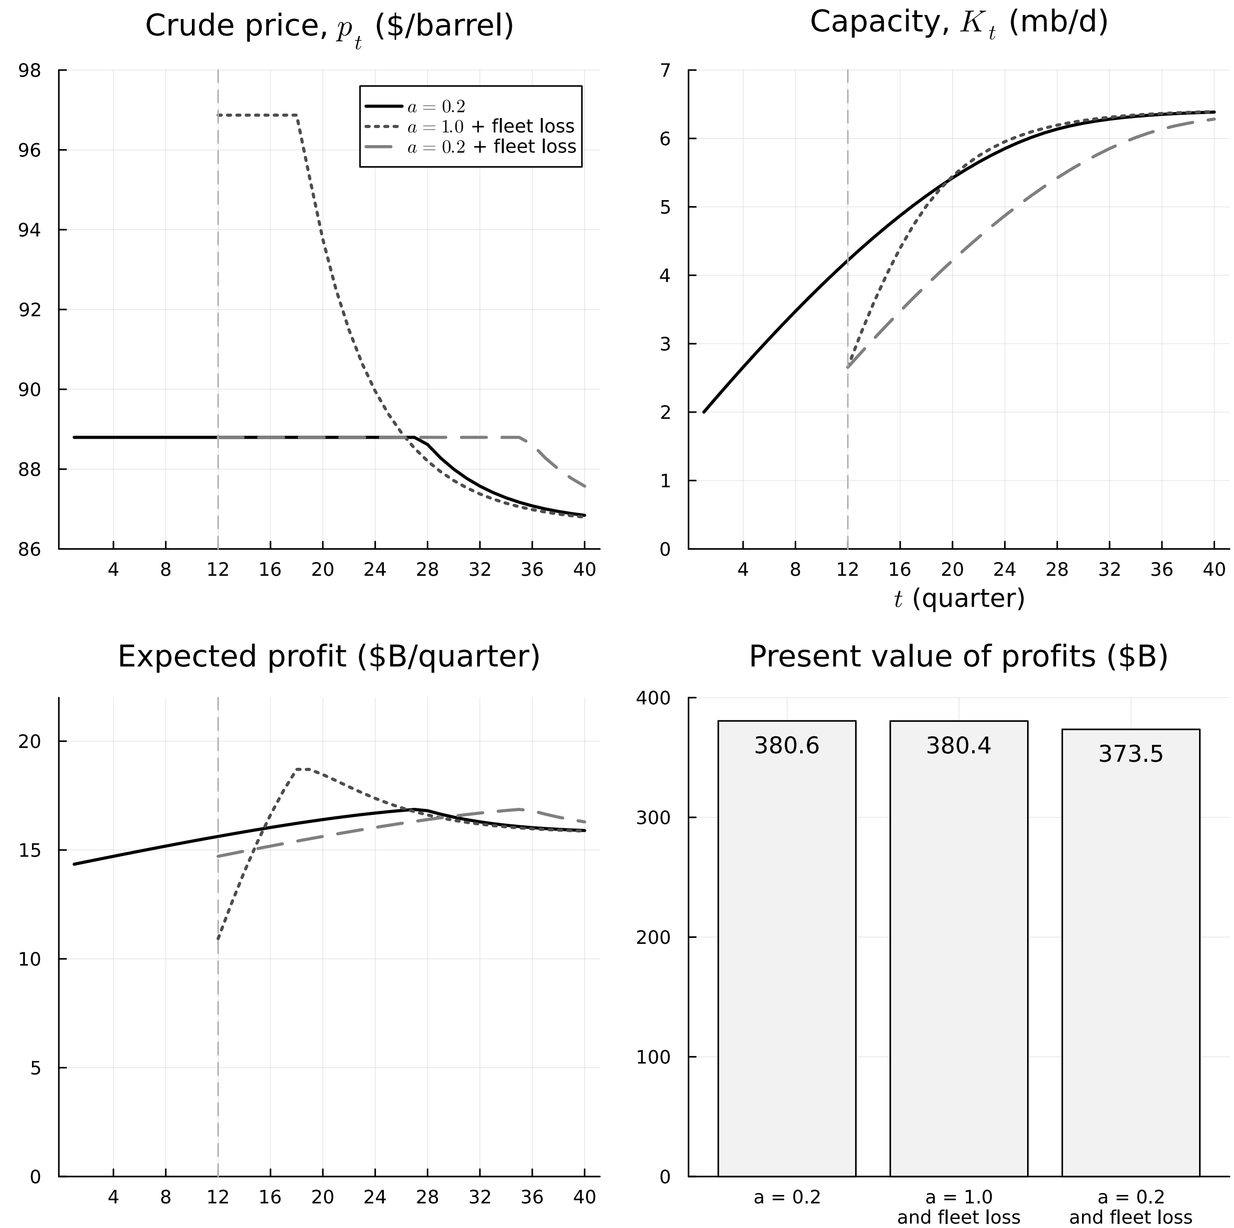

Figure C11: A comparison of the outcomes with and without unexpected changes after the 12^{th} quarter. The solid line represents the low enforcement case ($a = 0.2, \tau = 10$) throughout without shadow fleet loss. The dotted line represents the case with the destruction of the shadow fleet accumulated between the beginning of the 4^{th} and 12^{th} quarters and ramping up to perfect enforcement ($a = 1$) after the beginning of the 12^{th} quarter. The dashed line shows the same shadow fleet destruction but maintaining low enforcement. The present value of profits is calculated for post-targeting periods, discounted to $t = 12$. The line graphs display only the first 40 quarters of the 80-quarter simulation.

ously. 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 the present value of Russia's 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 the Russian post-targeting present value of expected profits imperceptibly (by .05%)—from \$380.6 billion to \$380.4 billion.

The dotted lines in Figure C11 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 C12 shows the results for the case of perfect enforcement of the cap before the 12^{th} quarter. The solid line reproduces the baseline results. The dotted lines represent the trajectories after the 12^{th} 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 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 post-targeting present value of Russian profits to increase by 1.2%.—from \$384.3 to \$389.0 billion.

The dashed lines in Figure C12 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 post-targeting 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

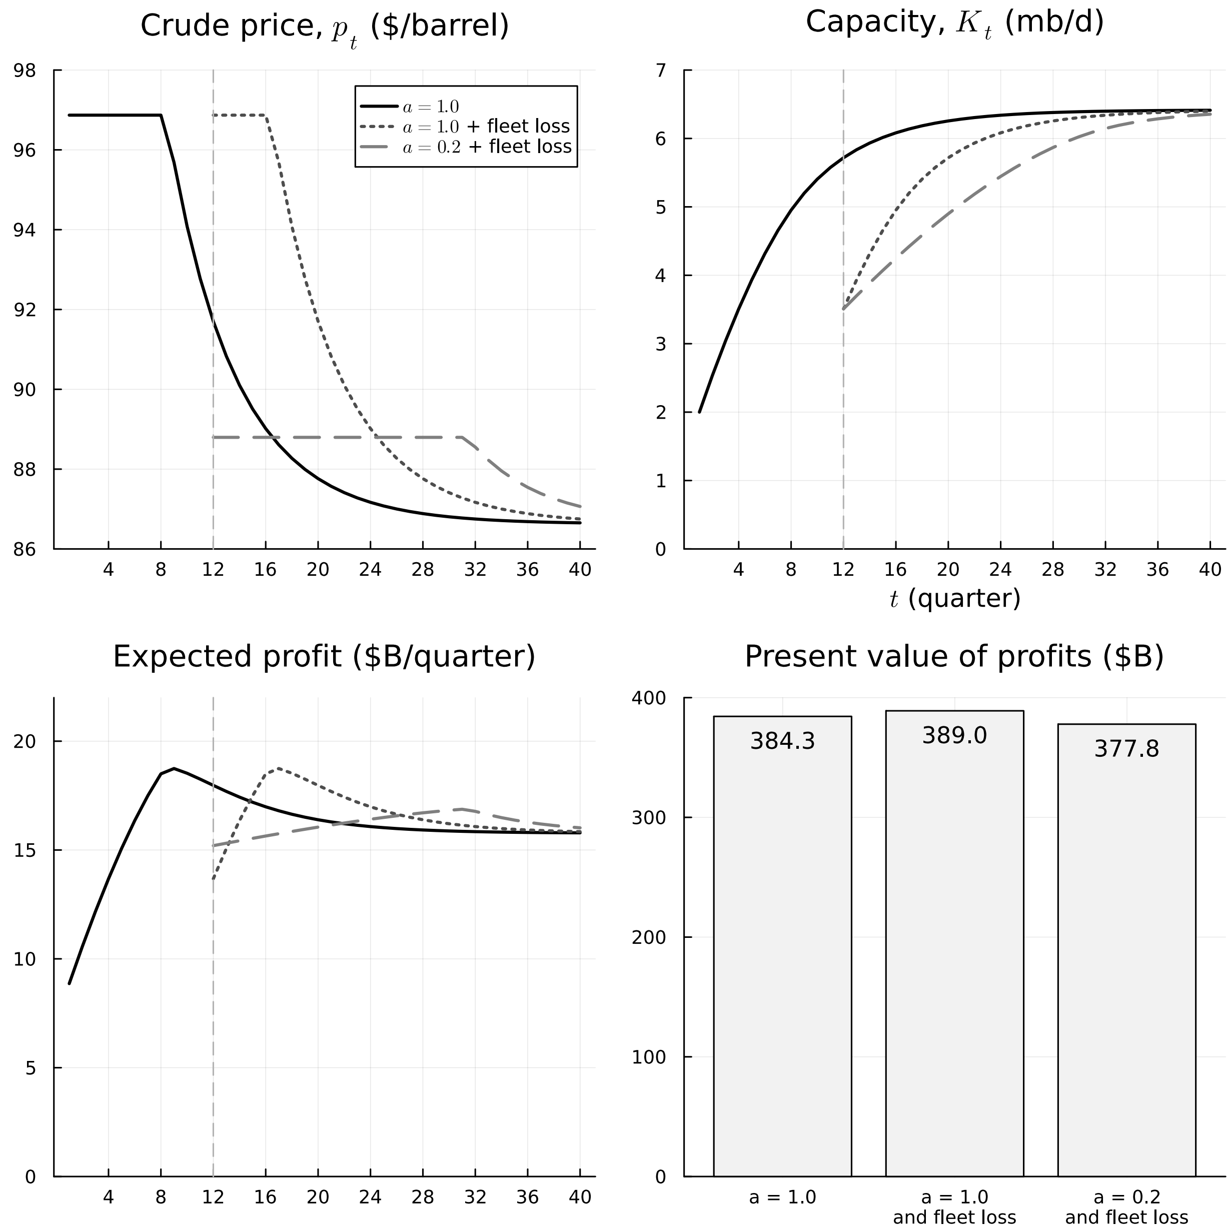

Figure C12: A comparison of the outcomes with and without unexpected changes after the 12^{th} quarter. The solid line represents the perfect enforcement case ($a = 1$) throughout without shadow fleet loss. The dotted line represents the case with the destruction of the shadow fleet accumulated between the beginning of the 4^{th} and 12^{th} 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 12^{th} quarter. The present value of profits is calculated for post-targeting periods, discounted to $t = 12$. The line graphs display only the first 40 quarters of the 80-quarter simulation.

the 12^{th} 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.[8] In the extreme case where fleet capacity was eliminated entirely, the post-targeting 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.

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

References

Baumeister, C., and Hamilton, J. D. (2019). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review*, *109*(5), 1873–1910.

Henderson, J., and Fattouh, B. (2016, February). Russia and OPEC: Uneasy partners. *Oxford Institute for Energy Studies*. Retrieved from https://www.oxfordenergy.org/wpcms/wp-content/uploads/2016/02/Russia-and-OPEC-Uneasy-Partners.pdf

International Energy Agency. (2022). *Oil Market Report - January 2022.* Paris. Retrieved from https://www.iea.org/reports/oil-market-report-january-2022

International Energy Agency. (2024a). *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 Energy Agency. (2024b). *Oil Market Report - January 2024.* Paris. Retrieved from https://www.iea.org/reports/oil-market-report-january-2024

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., and Murphy, D. P. (2014). The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, *29*(3), 454–478.

Kilian, L., Rapson, D., and Schipper, B. (2024). *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

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

Wachtmeister, H., Gars, J., and 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