NBER WORKING PAPER SERIES

MINING WITH ENVIRONMENTAL RISK

Christopher Costello
Charles D. Kolstad

Working Paper 21325
http://www.nber.org/papers/w21325

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2015, Revised September 2022

We thank Michael Hanemann, Dan Kaffine, Steve Salant, Marc Conte, and seminar participants at ASSA, UC Davis, University of Wyoming, and Oregon State University for helpful comments on previous versions of this paper. A previous version was circulated as "Mining and Quasi-Option Value." The authors declare no funding source and no conflicts of interest for this research. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Mining with Environmental Risk
Christopher Costello and Charles D. Kolstad
NBER Working Paper No. 21325
July 2015, Revised September 2022
JEL No. Q31,Q32,Q38,Q52

ABSTRACT

Environmental concerns dominate modern-day decisions about mining. While the minerals extracted are surely valuable, mining of natural gas, deep seabed minerals, rare earth metals, and traditional ore is often fraught with environmental uncertainty. We examine how this uncertainty affects the optimal decision of if, and when, to mine. When environmental damage from mining is known, the socially optimal timing depends straightforwardly on the magnitude of the damage relative to these damages in the rest of the world. But when environmental damage is uncertain, and its magnitude is learned over time, an option value arises, which fundamentally alters the mining decision. This decision depends on the costs and benefits of mining at different times, which are innately linked for non-renewable resources by Hotelling's rule. Using this insight, we find that any uncertainty over environmental costs can make it optimal to delay mining; this occurs even when expected environmental costs are low or even negative. We show conditions under which it is optimal to postpone the mining decision indefinitely, and conditions when it is optimal to postpone only for a finite duration. We use these insights to derive, for the first time, the equilibrium outcome of an entire industry of decentralized mine owners who all face an incentive to delay to acquire improved information. This gives rise to strikingly different price and extraction paths than are currently understood. One such outcome is that price paths flatten relative to what Hotelling theory predicts, consistent with empirical findings that have puzzled the literature.

Christopher Costello
Bren School of Environmental
Science & Management
University of California, Santa Barbara
Santa Barbara, CA 93106
and NBER
costello@bren.ucsb.edu

Charles D. Kolstad
Stanford University
366 Galvez Street (Room 226)
Stanford, CA 94305-6015
and NBER
ckolstad@stanford.edu

Mining with environmental risk*

Christopher Costello[†] and Charles D. Kolstad[‡]

September 12, 2022

Abstract

Environmental concerns dominate modern-day decisions about mining. While the minerals extracted are surely valuable, mining of natural gas, deep seabed minerals, rare earth metals, and traditional ore is often fraught with environmental uncertainty. We examine how this uncertainty affects the optimal decision of if, and when, to mine. When environmental damage from mining is known, the socially optimal timing depends straightforwardly on the magnitude of the damage relative to these damages in the rest of the world. But when environmental damage is uncertain, and its magnitude is learned over time, an option value arises, which fundamentally alters the mining decision. This decision depends on the costs and benefits of mining at different times, which are innately linked for non-renewable resources by Hotelling's rule. Using this insight, we find that *any* uncertainty over environmental costs can make it optimal to delay mining; this occurs even when expected environmental costs are low or even negative. We show conditions under which it is optimal to postpone the mining decision indefinitely, and conditions when it is optimal to postpone only for a finite duration. We use these insights to derive, for the first time, the equilibrium outcome of an entire industry of decentralized mine owners who all face an incentive to delay to acquire improved information. This gives rise to strikingly different price and extraction paths than are currently understood. One such outcome is that price paths flatten relative to what Hotelling theory predicts, consistent with empirical findings that have puzzled the literature.

JEL Classifications: Q3; H2; H4
Key words: Hotelling's rule; option value; quasi-option value; mining; environmental externalities

1 A Fable

Once upon a time, in a faraway land, a rich deposit of gold was found. It was so rich that the net profit from mining would be 50 billion Cubits (the local currency). A careful

*We thank Michael Hanemann, Dan Kaffine, Steve Salant, Marc Conte, and seminar participants at ASSA, UC Davis, University of Wyoming, and Oregon State University for helpful comments on previous versions of this paper. A previous version was NBER Working Paper 21325. The authors declare no funding source and no conflicts of interest for this research.

†4410 Bren Hall, UC Santa Barbara and NBER, costello@bren.ucsb.edu, corresponding author
‡Stanford University, UC Santa Barbara, RFF and NBER, ckolstad@ucsb.edu

environmental review revealed almost no environmental risk. There was a 99% chance of no harm at all, though the review did find a small (1%) chance of 51 billion Cubits in environmental damage. The true environmental damage would not be revealed for decades. The King, being a benevolent sort, wanted to appropriately balance the benefits to his kingdom against these environmental risks, so he appointed a three-member Council of Elders to advise him.

The first Advisor submitted his report that very day. He concluded that this was a good project for the kingdom, reasoning that the expected 0.5 billion Cubit environmental cost was a small price to pay for 50 billion Cubits in profit.

The second Advisor was not so sure. She reasoned that an option value must be considered: "If it turns out that environmental costs are indeed high, then the project cannot be 'undone', so this introduces a value of postponing the project." But after further consideration, she concluded that this option value was insufficient to delay the project, both because the benefits vastly exceed the expected costs and because the long time-frame for learning would shrink the present value of the project. She, too, advised proceeding with the project.

The third Advisor listened closely to these arguments and concluded that while they both contained a kernel of wisdom, they were both incomplete. She reasoned that the combination of an exhaustible resource with an option value introduced a new fundamental insight about this project and all projects like it. On the basis of her reasoning, she advised the King to delay.

In this paper, we will side with the third advisor and in so doing, will develop new insights on quasi-option value in the decision to mine an exhaustible resource. The punchline of the theory is that even in cases when cost-benefit analysis (Advisor #1) and standard quasi-option value alone (Advisor #2) suggest it is optimal to develop a natural resource, it is often economically advantageous to delay (Advisor #3).

2 Introduction

Two seminal results from the last century of environmental and resource economics are integrated in this paper, providing new insights on the timing of extraction of exhaustible resources and informing important contemporary policy debates. One is the identification of quasi-option value as potentially important for the preservation of natural environments (Arrow and Fisher 1974; Henry 1974).[1] The concept arose in part out of a proposal to dam the Snake River in Idaho, which would have irreversibly destroyed many aspects of the natural environment with uncertain environmental consequences. Cost-benefit analysis, even taking uncertainty into account, suggested damming the river yielded positive expected net benefits. Yet absent from the standard cost-benefit analysis was the fact that damming the river would foreclose on any possibility of future enjoyment of it. What was missing was a cost associated with taking an irreversible action that eliminates the flexibility to act conditional on new information; this is the quasi-option value.

[1] The term 'quasi-option value' was used by Arrow and Fisher (1974) to distinguish it from 'option value,' introduced by Weisbrod (1964) to capture the effect of uncertain demand for natural environments (an apparently closely related concept).

This insight now pervades theory and practice. Any given project (such as a dam) has expected costs (construction, engineering, environmental) and benefits (flood protection, water and energy provision). When costs or benefits are uncertain and information about them will be revealed over time, the quasi-option value quantifies the value of postponing the development decision and thus preserving the option of using the (undammed) watershed in the future. Crucially, though, this does not necessarily imply that the project should be postponed. Rather, it *nudges* the project toward postponement. But if expected benefits are significantly higher than expected costs, the quasi-option value will be immaterial, and the dam will optimally be built anyway. In this regard, we may surmise that quasi-option value will have practical significance only in a narrow set of circumstances (those in which the benefits of development slightly outweigh the expected costs).

A second seminal result of the last century of environmental economics is the theory put forward by Harold Hotelling regarding the dynamics of prices for exhaustible resources. As robustly demonstrated by Hotelling (1931), efficient market-based production of non-renewable resources involves a temporal arbitrage condition (rents rise at the rate of discounting) that inextricably links benefits and costs of a mining project over time. One implication of this result is a sort of indifference about when to mine; the increase in prices over time exactly offset the cost of delay from discounting. This implication seems to suggest there is little or no cost to delaying mining until improved environmental information becomes available. Importantly, no such intertemporal indifference property exists for other social decisions, such as whether to dam a river.

We merge the theories of Hotelling and quasi-option value to address the question of whether, and when, to engage in environmentally-risky mining activity. Applying the structural restrictions from Hotelling theory to the general theory of quasi-option value sheds new light on this contemporary debate and gives rise to significant policy implications. By exploiting both of these economic theories simultaneously, we are able to develop new insights on the "when to mine" question and in so doing, we derive a new quasi-option value for exhaustible resources. This class of problems is of tremendous current policy relevance for issues such as deep seabed mining, rare earth mineral extraction, and hydraulic fracturing for gas and oil, all of which are fraught with environmental risk.

We find that this problem is fundamentally different than, say, the decision about whether to build a dam. Instead, because the resources we consider are exhaustible, Hotelling theory largely neutralizes the role of time (because rent goes up at the discount rate) whereas quasi-option value exploits the tension between time (the discount rate favors developing now) and information acquisition (develop later). This linkage turns out to be crucial vis-à-vis the role of quasi-option value for exhaustible resource projects. For this class of problems, the resource in-situ is expected to increase in value over time, according to the Hotelling rule; this is not generally true for canonical applications of quasi-option value such as environmental values threatened by a dam or other development. Remarkably, exhaustible resources have played almost no role in the quasi-option value literature.[2]

[2] Traeger (2014) provides a comprehensive treatment of quasi-option value in environmental settings, but does not analyze the problem for non-renewable resources. Hoel (1978) sets up a problem similar to what is found in this paper − uncertain extraction costs with uncertainty removed at some known future time. He focuses on risk aversion generating over-extraction and suggests a tax to support efficient extraction. He does not discuss quasi-option value nor the difference between efficient and inefficient decision makers as is

At the market level, the Hotelling rule implies indifference between development now and in the future for the marginal mine. This intertemporal indifference suggests that even a small quasi-option value may be enough to tip the optimal development decision of a marginal mine towards significant delay. In contrast, for a standard resource project (e.g. dam construction), there is a real loss from delay, in terms of deferred but otherwise constant net benefits (absent exogenous assumptions about how net benefits might change over time); in that setting a small quasi-option value would not be expected to change the development decision. The main point of this paper is that exhaustible resources are fundamentally different from all other market goods in the context of information and irreversibilities. The interaction between the Hotelling principle and quasi-option value can generate fundamentally different results and insights on the development of an exhaustible resource with uncertain, but ultimately knowable, benefits of preservation.

As with many theoretical inquiries in economics, our problem is motivated by a real-world policy dilemma – the potential development of the Pebble Gold Mine in Alaska (Parker et al. 2008; Holley and Mitcham 2016). The deposit is one of the largest and richest deposits of gold, copper, and molybedenum in the world. The area is also home to a globally significant salmon fishery and other ecological and cultural resources. Early assessment of the risks of mining to fishery and ecosystems were considered significant enough for the US EPA to invoke the Clean Water Act to substantially restrict the development of the mine (Environmental Protection Agency 2014), which seemed to doom the project. But in mid-2017 the D.J. Trump administration proposed removing EPA objections (Dennis 2017), which would have breathed new life into the mining project. Somewhat ironically, on August 4, 2020, Donald Trump Jr. tweeted that "The headwaters of Bristol Bay and the surrounding fishery are too unique and fragile to take any chances with." This is a plain English variant on the theme of this paper. Shortly thereafter, the US Army Corps of Engineers appeared to place additional roadblocks to the project. Then Mr. Trump lost the election to now President Biden who had pledged to cancel the Pebble Mine. In September of 2021, under the Biden Presidency, the EPA announced it would take action under the Clean Water Act to overturn Trump's actions to facilitate development of the mine (Grandoni and Partlow 2021). Clearly the issue is controversial. The project is an excellent example of a mine that appears to be privately profitable but with potentially significant and highly uncertain social (environmental) costs. See also Holley and Mitcham (2016) for a discussion of the social license to operate and Narula (2014) for a perspective on the politics of the project.

Although we do not take a position on the Pebble Mine, this paper does concern decisions by public officials about whether to proceed with immediate development of an exhaustible resource or to postpone development until uncertain environmental costs are better known - the situation Alaskan and US EPA officials have found themselves in for the case of the Pebble Mine. Other equally-compelling contemporary examples include whether to proceed with hydraulic fracturing ("fracking") for natural gas, logging old-growth forests, or engaging in deep seabed mining for copper, gold, or rare earth metals. Each of these examples concerns real-world policy questions that pit the net private benefits of extraction against the uncertain environmental costs of mining. One contribution of this paper is to show that a simple cost benefit analysis, which compares expected benefits to expected costs, will often get the wrong

done in this paper.

answer. A key insight that arises from the theory is that deciding to postpone mining until better information becomes available forfeits none of the resource underground, and indeed often forfeits none of its value, since its value is rising over time, as per Hotelling.

For the remainder of this paper, we remain agnostic about the particular application. We are concerned with a mine that has certain private costs and benefits, but uncertain external social costs (we think of these as environmental, but they can take many other forms). A decision must be made regarding development of the mine, encompassing private and social costs. The mine has an owner, who is concerned only with private costs and benefits; the mine owner will be assumed to be completely dynamically rational from an economic point of view, which of course, among other things, means the mine will ignore external costs. The social planner is a mine agency that oversees the mine and accounts for all private and social costs and benefits. The mine agency may impose costs on the mine owner that effectively internalize social costs or the mine agency may evaluate the overall costs and benefits and regulate the output from the mine over time, including deciding whether to allow the mine to proceed (or more generally, how much if anything to produce at any point in time).

Thus the problem we consider is the case of a planner (the mine agency) making a decision on whether to allow a mine to proceed in a context where environmental costs are uncertain.[3] For example, these uncertain environmental costs could reflect the possibility of drinking water contamination (from fracking), salmon river destruction (from gold mining), or biodiversity loss (from deep seabed mining). We allow for the possibility that environmental uncertainty can be resolved at some point in the future. To highlight the potency of quasi option value in this setting, we follow Traeger (2014) and explicitly account for the level of sophistication of the planner. We contrast the case of a naïve planner, who ignores the fact that information will be obtained in the future, with a fully rational planner who utilizes all available information, making dynamically efficient decisions (a socially efficient planner). The comparison between a naïve and a rational planner maps well to the Pebble Mine example that motivated this work, with regulatory authorities making a determination on the social desirability of allowing the mine to proceed.

Our analysis relies fundamentally on the concept of quasi-option value, which has its roots in capital theory with irreversible investment (Arrow 1968; Arrow and Kurz 1970). Under uncertainty about the value of the natural environment, if uncertainty is reduced over time, there is a value associated with postponing irreversible development, a cost of development that should be reflected in decision-making about whether to develop immediately or not. Over the years there has been some confusion in the literature about the distinction between option value and quasi-option value. Grappling with the theoretical distinction is muddied further by the fact that financial economics has yet another concept of option value, developed for application to natural resource economics by Dixit and Pindyck (1994). Fortunately, a number of authors, particularly Hanemann (1989), Mensink and Requate (2005), and Traeger (2014), have provided unifying syntheses of and clarity among these three distinct concepts. Furthermore, terminology appears to have evolved to distinguish between the 'Arrow-Fisher-Henry' quasi-option value and the 'Dixit-Pindyck' option value. For clarity in this paper, we will use the term 'Arrow-Fisher-Henry quasi-option value' interchangeably

[3]While we have framed the problem as a decision maker who is uncertain about environmental cost, the story applies equally to any uncertainty over extraction costs that will be revealed over time.

with ‘quasi-option value.’

Uncertainty is of fundamental importance to our results. But uncertainty alone is insufficient; we also require that uncertainty can be resolved, at least in part, over time. As clearly pointed out by Hanemann (1989), quasi-option value is the value of the information ultimately received, conditional on postponing irreversible development. Thus quasi-option value is related to the value of information, and, importantly, does not require risk aversion.[4]

Much of our analysis focuses on the behavior of a regulatory body concerned with external costs and a marginal mine owner who takes exhaustible resource prices as given, in an atmosphere of uncertainty and learning over time. The affirmative decisions of both the regulator and the mine owner are necessary for mining to proceed. Our main contribution is to show that strong incentives exist to delay mining until uncertainty is resolved.

In the next section we develop a simple model of the timing of extraction of an exhaustible resource with uncertain costs, where uncertainty can be resolved over time. We contrast the decision made by the naïve planner (who ignores learning about environmental costs over time) with the decision made by a sophisticated (or socially efficient) planner (who completely accounts for the possibility of new information over time). In Section 4 we seek a decentralized mechanism (a corrective tax on extraction) that can decentralize the decision-process of the sophisticated planner; we also provide a numerical illustrative example to show how the dynamics play out. As an extension, we then derive the equilibrium behavior of a competitive industry in which all mine owners face an option value in Section 5. We conclude in Section 6.

3 Model of Social Decisions

We initially consider a single mine owned by a small producer (e.g., an Alaskan mine), denoted with subscript A, operating as a price-taker within a global market for the non-renewable resource produced by the mine. To keep things simple, assume mine A contains a single unit of the resource and A’s problem is to determine when to deplete her mine (extraction/depletion occurs all at once when it does occur). Extraction must be profitable for the mine owner of course but the mine is also subject to oversight by a regulator/planner who will take into account uncertain environmental costs, ε_A, and will permit mining only if it is in society’s interest to do so. The mine owner does not factor external environmental costs directly into her decisions.

To extract the resource, mine owner A faces known private extraction cost γ_A. The world

[4] Traeger (2014) provides a clear synthesis of option and quasi-option value. He sets up a simple two period model with learning and distinguishes three types of decision-makers regarding irreversible development: (1) a sophisticated decision-maker (s) who anticipates learning and may postpone part or all of her decision until after information is acquired; (2) a less sophisticated decision-maker (p) who makes all decisions ex ante, before uncertainty is resolved, but allows the possibility to develop in either period; and (3) a pure naïve decision maker (n) who makes an all or nothing decision in the first period based on expectations (depending on how the problem is set up, n and p may be equivalent). The difference between the value of development for s vs. p is the quasi-option value; the difference for p vs n is the simple option value (unrelated to learning); and the difference for s vs. n is the full value of sophistication (which may be greater than the quasi-option value, depending on whether the simple option value is non-zero, which depends on the nature of payoffs and uncertainty).

price at time t is $P(t)$, the interest rate is r, the initial period is $t = 0$, and extraction of a unit of the resource in the rest of the world entails cost c_w (this is the cost the market sees, which may be the private cost of extraction or a social cost with internalized external costs). Assuming a competitive global market with fixed reserves, the price path follows Hotelling's rule over time:

$$P(t) = e^{rt}P(0) - c_w(e^{rt} - 1) \tag{1}$$

If $c_w = 0$, then the global price of the resource rises at the rate of interest, r. But if $c_w > 0$, then rent rises at the rate of interest and thus the global price of the resource rises more slowly than the rate of interest.

Equation 1, together with total initial global reserves R and global demand (assumed to have a choke price $\tilde{p} > c_w$), will determine the initial price, $P(0)$, the price path, $P(t)$, and the exhaustion date T, all of which are exogenous to owner A. More complex models of mining conditions and behavior are considered later in this paper, but do not qualitatively affect our results.

3.1 Mine-owner's private decision

We begin by asking: When would owner A like to extract her resource? Since the price path $P(t)$ is exogenous to owner A, she can take it as given. If she decides to mine at date s, her present value profit is:

$$\pi_A(s) = e^{-rs}\left[P(s) - \gamma_A\right] = P(0) - c_w + e^{-rs}(c_w - \gamma_A) \tag{2}$$

This function is increasing in s if A's costs are high ($\gamma_A > c_w$) and decreasing in s if A's costs are low ($\gamma_A < c_w$). So the privately optimal mining-time for owner A, t^*_A, is given by:

$$t^*_A = \begin{cases} 0 & \text{if } \gamma_A < c_w \\ \text{any } t \in [0, T] & \text{if } \gamma_A = c_w \\ T & \text{if } c_w < \gamma_A \leq \tilde{p}) \\ \text{Never} & \text{if } \gamma_A > \tilde{p} \end{cases} \tag{3}$$

This simple result has a strong economic intuition: When the private cost in region A is small relative to the rest of the world, it pays to extract immediately because the rent in region A will rise more slowly than at the rate of interest. When the private cost in region A is large but less than the choke price, the opposite holds: although extraction would be profitable, rent is rising faster than the rate of interest so it is optimal to defer extraction until all other mines are exhausted, which occurs at time T. And when the private cost is sufficiently large, so that even the highest price possible ($\tilde{p}$) would not justify the cost, the resource should be left in the ground indefinitely.

3.2 The social planner's public decisions

The mine owner's problem studied above assumes that owner A is aware of her private costs and benefits and makes her mining decisions in an economically rational manner. But the mine A is overseen by a planner, representing society; the planner must account for all social costs. Thus, the true social cost of mining in region A is given by the sum: $c_A = \gamma_A + \varepsilon_A$. Since the planner will attempt to capture both (certain) private and (uncertain) public environmental costs, she will view c_A as a random variable, with mean denoted by $\bar{c}_A$:

$$\bar{c}_A = \gamma_A + E[\varepsilon_A] \tag{4}$$

The variable $\bar{c}_A$ can thus be thought of as the *expected social cost* of mining in region A.

The timing of receipt of information and its effect on uncertainty are also important. We assume that at time 0 the environmental cost of mining is uncertain, so c_A is uncertain and its probability distribution is given by $f(c_A)$. Following the quasi-option value literature, we assume information revealing the true value of environmental cost (ε_A), and thus social cost, c_A, will become available at some known point in the future.

As mentioned earlier, we consider two plausible models for the planner to make decisions about proceeding with the mine. This is in the spirit of the Arrow and Fisher (1974) model of regulator behavior. One representation of planner behavior is fully dynamically rational, taking into account the fact that information will be obtained in the future. We term this the rational planner. The other model of planner behavior is what we call a naïve planner, who is well-aware that there will be environmental costs of mining, and will attempt to integrate those into mining decisions, but she makes her timing decision ignoring the fact that information will be acquired in the future. The reason for including the naïve planner is that it is consistent with the literature on option value and consistent with the practice of many actual cost benefit analyses. In fact, most actual cost benefit analyses in government ignore uncertainty, let alone the fact that uncertainty may evolve over time, with information being revealed in the process. We calculate the optimal time to mine (and associated expected payoff) for both types of decision makers.

The Hotelling rule implies that the price path for the nonrenewable resource is closely tied to the extraction cost faced in the rest of the world (c_w), as in Equation 1. Thus, the timing of mining in A will hinge on the comparison of c_A with c_w. Perhaps the most interesting and empirically-relevant case is when expected cost in region A is equal to the cost faced by other mine owners in the world (so $\bar{c}_A = c_w$). After all, if region A's cost were much lower, we would have expected region A to have already exploited the mine. And if region A's cost was much higher, postponing mining would be an obvious choice in region A. Thus, we regard the condition $\bar{c}_A = c_w$ (approximately) as the typical state of affairs, and we will adopt it as the benchmark for much of our analysis. In this case, we will show that the naïve planner is completely indifferent about when to mine over the interval $[0, T]$. That is, the naïve planner in A is just as happy to mine today, next year, or in a decade. This linking of current rents with future rents is a feature of all non-renewable resource extraction problems, but is absent from all traditional quasi-option value models, and is what allows us to derive new insights about mining and quasi-option value.

The naïve decision maker sees uncertain cost $c_A = \gamma_A + \varepsilon_A$, but ignores new information so the distribution over A will remain $f(c_A)$ forever.[5] In this case, decision maker A's expected

[5]Equivalently, we could assume that decision maker A learns the true value of A only after mining has occurred.

payoff from a decision to mine at date s is given by:

$$E[\pi(s)] = \int_{-\infty}^{\infty} \Big(P(0) - c_w + (c_w - c_A)e^{-rs}\Big) f(c_A) dc_A \tag{5}$$

$$= P(0) - c_w + e^{-rs}(c_w - \bar{c}_A) \tag{6}$$

Equation 6 bears a striking resemblance to the payoff achieved by the private mine owner herself in Equation 2. The key difference is that the naïve decision maker prices in the *expected* externality cost (because $\bar{c}_A = \gamma_A + E[\varepsilon_A]$). But it turns out that uncertainty plays only a trivial role - the optimal mining decision depends on the expected value of c_A, but not on any other features of the distribution of c_A. The optimal mining-time for naïve decision maker A is given by:

$$t_A^* = \begin{cases} 0 & \text{if } \bar{c}_A < c_w \\ \text{any } t \in [0, T] & \text{if } \bar{c}_A = c_w \\ T & \text{if } c_w < \bar{c}_A \leq \tilde{p}) \\ \text{Never} & \text{if } \bar{c}_A > \tilde{p} \end{cases} \tag{7}$$

So the optimal timing desired by the naïve decision maker requires that expected environmental costs are priced in. Doing so replaces the private cost (γ_A) with the expected social cost ($\bar{c}_A = \gamma_A + E[\varepsilon_A]$) and proceeds accordingly. Inspecting Equations 3 and 7, adding $E[\varepsilon_A] > 0$ will tend to push back the mining decision and seems to accord with how many real-world policy makers think about mining decisions on public lands: Do the expected benefits outweigh the expected costs?

While the naïve decision maker thinks she is doing the socially-desirable thing by pricing in the environmental cost of mining in expectation (Equation 7), this approach is still incomplete because it ignores potentially useful information. To model a sophisticated decision maker, we now add learning to the optimal mining timing under uncertainty. While many forms of learning are possible, we adopt a simple version where sophisticated planner in A learns the true environmental cost either upon mining or at a known future date $\tau < T$, whichever occurs first. That she would learn her costs upon mining is not too far-fetched. For example, tailings from the Pebble Mine either do, or do not, compromise salmon populations and fracking either does, or does not, pollute drinking water; these will be revealed after mining has taken place. If she chooses to defer mining until after date τ, we assume that the environmental cost will be revealed, with certainty, at date τ. This is meant as a heuristic that captures the idea that over time, exogenous information may be revealed that would help decision maker A identify the true environmental cost. Again, following the environmental example, scientific information may accrue over time that narrows A's uncertainty over the true external cost of mining.[6]

The prospect of learning the true environmental cost (ε_A), and thus the true social cost (c_A) prior to mining is an enticing one. If deferring the mining decision is not too costly, and provided that A might learn something useful (i.e. something that might change her optimal mine-time), then the planner may wish to defer, at least until date τ. Whether this

[6]To convince yourself of the plausibility of this claim, consider the dramatic improvement in environmental information on mining or old growth logging that has accumulated over the past 50 years.

is the case depends on an interplay among (1) the *expected* cost, $\bar{c}_A$, (2) others' costs, c_w, (3) the date of information revelation τ, and (4) the shape of distribution $f(c_A)$. In this section we solve for this relationship and highlight its key features that affect the optimal timing decision. The optimal timing decision for a sophisticated decision maker turns out to hinge on how A's expected cost compares to the actual cost of others ($\bar{c}_A \lesseqgtr c_w$) and on the choke price. We examine each case in turn.

3.2.1 Rational Planner and "Typical" expected social cost: $\bar{c}_A = c_w$

We begin with the most natural starting point in which the social cost of mining in region A is expected to be similar to others' costs, so $\bar{c}_A = c_w$. Inspecting the result in Equation 7, it is tempting to think that if A's cost is expected to be the same as everyone else's costs, then A would be indifferent about when to mine. This initial logic is incorrect, which leads to our main result:

Proposition 1. *For any positive amount of uncertainty in social cost, c_A, if $\bar{c}_A = c_w$, then it is always optimal for a dynamically rational planner to postpone mining until at least date τ.*

Proof. Formal proofs reside in the Appendix. □

While the formal proof requires detailed analysis, the intuition behind this result is straightforward. The basic idea is to compare the expected social payoff from mining prior to τ (which we call V_1) with the expected social payoff from postponing until at least date τ (which we call V_2). Committing to mining prior to τ, gives the expected payoff of the naïve planner (because no new information becomes available prior to τ). But the naïve planner obtains the same expected payoff regardless of when she mines (see Equation 7). So V_1 is identical to the expected payoff obtained when the planner *commits to mining at date* τ. And V_2 is the expected payoff from waiting until date τ and then deciding when (and whether) to mine. When dissected in this way, it is obviously better to maintain flexibility; i.e. to delay mining until at least date τ.

The flexibility is valuable for two separate reasons. First, if c_A turns out to be large, but not too large (this is revealed at date τ), then it will be optimal to mine at $T > \tau$, in which case $V_2 > V_1$. Second, if c_A turns out to be very large, then it will be optimal to never mine, in which case $V_2 > V_1$. In both of these cases, it is optimal to postpone until at least date τ. The third possibility is that it is revealed at date τ that costs are low (in which case it would have been optimal to mine at date 0). But *ex ante*, committing to mine at date 0 has a payoff identical to committing to mine at date τ, so the two payoffs are equal in that case. Taken together, when $\bar{c}_A = c_w$ it is always optimal to postpone the mining decision until at least date τ. This simple, yet dramatic result implies that when owner A's expected private cost of mining is "typical," uncertainty over the environmental cost of mining will always persuade the sophisticated planner to delay mining.

3.2.2 Rational planner and "High" expected social cost: $\bar{c}_A > c_w$

If instead the expected social cost of mining in region A is higher than c_w, then the result of the previous section is strengthened further. As expected cost rises, there is an incentive to push those costs farther into the future, which implies mining at a later date. Indeed, even a naïve planner chooses to postpone mining (until at least date τ) in this case (see Equation 7). The result is summarized as follows:

Proposition 2. *For uncertain social cost, c_A, if $\bar{c}_A > c_w$, then it is* always *optimal for the dynamically rational planner to postpone mining until at least date τ.*

The logic underlying the proof is similar to that in Proposition 1. Here, the *ex ante* expected payoff from mining prior to τ is lower than the expected payoff from committing to mine at date τ. And the expected payoff from postponing the decision until date τ is strictly greater than the expected payoff of committing to mine at date τ. So postponing the decision until date τ is always preferred to mining prior to τ.

3.2.3 Rational planner and "Low" expected social cost: $\bar{c}_A < c_w$

Propositions 1 and 2 reveal that provided A's expected social cost is equal to or greater than the global costs, it is always optimal to postpone the mining decision until better information arrives. But what if A's expected cost is lower than others' costs? In that case, an interesting tension arises. When $\bar{c}_A$ is low, there is an incentive for the naïve planner to allow immediate extraction (see Equation 7); this implies that delaying until τ would require her to sacrifice some expected returns. Whether this force is sufficient to overcome the benefit to a rational planner of preserving flexibility turns out to depend in interesting ways on features of the problem. Our main result here is that even though a simple naïve cost benefit analysis suggests that it is optimal to mine immediately, it is often optimal, instead, to postpone mining until at least date τ. This result is summarized below:

Proposition 3. *For uncertain social cost, c_A, if $\bar{c}_A < c_w$, then it may be optimal for the dynamically rational planner to postpone mining until at least date τ.*

The proof relies on carefully disentangling the tension between two factors and leveraging the observation that the rational planner would *never* want to mine in the open interval $(0, \tau)$. On the one hand, the fact that expected cost is low suggests that mining should commence immediately because the rent generated is more valuable as a standard investment (earning rate r) than it is in the ground. But on the other hand, there is a benefit to postponing the decision until at least date τ when the true cost c_A is revealed. Under this setup there is no benefit to postponing to some date $s < \tau$ because nothing can be learned, and the ex ante expected payoff is lower than it would be if mining took place at time 0. The benefit of postponing mining until at least τ should be clear: the true cost c_A might be relatively high, in which case it would be optimal to mine at date T, or it might be very high, in which case it will be optimal to never mine the resource. Neither of these options is available if A mines at date 0.

While Proposition 3 shows that it *might* be optimal to postpone mining until at least date τ, we would like to shed light on the factors that make this result more likely. This result is summarized as follows:

Corollary 1. *Under the conditions of Proposition 3, the payoff from postponing the mining decision (until τ) is increasing: (a) as τ becomes smaller and (b) with increasing uncertainty over $f(c_A)$.*[7]

These both accord with straightforward economic intuition.[8] As the learning time (τ) is smaller, so is the cost of delay, so the relative payoff from postponing gets larger. And as the distribution over social cost ($f(c_A)$) becomes more spread, the possibility of learning something very useful is increased, so this also increases the value of delay. Proposition 3 is one of the main results of this paper: for most expected social costs, the naive and rational planners behave the same; but for a range of expected social costs, the naive planner would approve mining whereas the rational planner would not.

3.3 Alternative Market Specifications

To derive the behavior of a single firm under uncertainty, we assumed a relatively simple market in which all mine owners (with the possible exception of owner A) share the same marginal cost of extraction. There are other competitive assumptions that could be adopted regarding the overall resource market: existence of a backstop technology; non-existence of a choke price; heterogeneous world market consisting of mines with constant but differing marginal extraction costs and reserves; and production with uncertain aggregate reserves and reserve additions occurring simultaneously with production; production and exploration with uncertainty about future reserves and/or future demand (Pindyck 1980).

Our results are qualitatively robust to these, and other, alternative competitive market specifications. All of these models result in a global price path, typically with rents rising at the rate of interest. Now introduce a small mine owner A who holds a single unit of the resource and can produce it for cost c_A. If c_A is known, A can perfectly "slot" herself into the queue - there is an optimal time t_A^* at which A should mine in order to maximize return. Suppose instead that c_A is uncertain, has the same mean as the cost that gave rise to t_A^*, and will be revealed at date $\tau > t_A^*$. It is straightforward to show in this model (in a manner similar to the analysis of Proposition 3 and Corollary 1 above) that conditions exist under which A will want to postpone the decision until date τ. It is also straightforward to show that owner A will never postpone to a date in the open interval: (t_A^*, τ), nor will she ever mine at a date prior to t_A^*. Thus, in a manner similar to the dynamics above, adding uncertainty to c_A may cause owner A to defer mining to a later date, or possibly defer it indefinitely.

The existence of a backstop technology is equivalent to demand having a choke price, at which demand drops to zero. Without a choke price, the terminal time T may become infinite but that does not change the analysis of the quasi-option value.

Deriving the equilibrium outcome when *all* firms face uncertainty is much more challenging, and reflects a much richer set of incentives and equilibrium conditions. We return to this case in Section 5.

[7] We induce "increasing uncertainty" by following the definition of "increasing risk" (though not necessarily "increasing variance") by Rothschild and Stiglitz (1970).

[8] And will be illustrated in a numerical example later.

4 Decentralizing the Mining Decision

The analysis thus far uses the temporal arbitrage result of Hotelling (1931) to show that it may be the case that a dynamically rational social planner would find it optimal to delay mining a non-renewable resource until the environmental consequences of doing so have been revealed. Yet a naïve planner, who prices in expected environmental cost, but ignores the possibility of learning about those costs and the tendency of natural resource rents to increase, may wish to permit immediate mining.

Rather than viewing the planner's job as accepting or rejecting the mine, we examine how a sophisticated planner could decentralize its optimal decision by imposing a tax on the mine to internalize the externalities (both environmental and informational). Both the naïve and sophisticated planner will internalize the expected environmental costs trivially by imposing a tax equal to $E[\varepsilon_A]$. But internalizing the informational costs (quasi-option value) associated with learning will require a more nuanced approach.

Hanemann (1989) re-formulates the quasi-option value model in an intriguing manner that will become useful here. He shows that the quasi-option value can be thought of as an additional "tax" on development (in addition to the expected environmental costs). Borrowing from Hanemann's insights, the purpose of this section is to add to expected environmental costs an additional informational tax, which we will denote by Q. The two part tax, imposed if and only if the owner chooses to extract prior to τ (a penalty for "early withdrawal"), will give rise to efficient (socially optimal) mining activity, fully decentralizing the sophisticated planner's decisions.

It is important to underscore that there are two parts to the tax. One is to correct the environmental externality. The other is to correct the informational externality. Thus, the tax we analyze here is designed to isolate the quasi-option value component of optimal delay–what could be considered an informational externality since it is information about social costs that may be acquired.

Suppose the private mine owner analyzed above faces a tax, $E[\varepsilon_A]$, that must be paid whenever she mines, and an additional tax, Q, that must be paid in order to mine prior to date τ. The goal is to derive the value of Q that, if included as an additional private cost of mining prior to τ, would give rise to the exact same mine timing that would be chosen by a rational planner. Thus, we must carefully attend to the incentives of the private mine owner. In order to induce a delay in mining, when she would otherwise find it desirable to mine, we must impose a sufficiently high penalty on premature mining. On the other hand, this penalty cannot be too large, or she will delay even when it is socially optimal to mine immediately.

Let $V^*(\tau)$ be the present social value (i.e., at date 0) of the deposit, to a naïve planner who commits to *postponing the mining decision until at least date* τ, and thus $V^*(0)$ denotes the social value to the naïve planner of mining at date 0. We are interested in whether the naïve planner will approve mining at date 0 or delay until at least date τ. This decision simply depends on $V^*(0) \lesseqgtr V^*(\tau)$. Thus, if we *always* wanted to encourage mining delay, imposing a tax on extraction prior to τ of at least $V^*(0) - V^*(\tau)$ would do the job.

However, we do not *always* want her to delay mining. We only want her to delay mining in cases in which the dynamically rational planner would delay mining. To analyze what society would like to induce the mine owner to do requires analyzing the rational planner's problem. Let $\hat{V}(\tau)$ be the expected present social value return, to a rational planner, who commits to *postponing until at least date* τ, and let $\hat{V}(0)$ denote the expected value of mining at date 0. Then, the socially efficient decision about whether to mine at date 0 or delay until at least τ will depend on $\hat{V}(0) \lesseqgtr \hat{V}(\tau)$. We only want the private mine owner to delay mining in cases in which $\hat{V}(0) < \hat{V}(\tau)$. Making use of the insights from Hanemann (1989), define the following informational tax:

$$Q = (V^*(0) - V^*(\tau)) - \left(\hat{V}(0) - \hat{V}(\tau)\right) \tag{8}$$

If this tax (plus expected environmental costs) is imposed on the private mine owner should she mine prior to τ, then such an owner would compare $V^*(0) - Q$ and $V^*(\tau)$, and would postpone to date τ if and only if $V^*(0) - Q < V^*(\tau)$, which is equivalent to

$$V^*(\tau) + \hat{V}(0) - \hat{V}(\tau) < V^*(\tau) \tag{9}$$

Simplifying reveals that the decision boils down to postponing mining if $\hat{V}(0) < \hat{V}(\tau)$, precisely the condition we had hoped to replicate. Thus the tax in Equation 8 "corrects" private owner's recalcitrance vis-à-vis new environmental information in the future. Imposing the tax Q on the private mine owner if she mines before τ will precisely align incentives - after incorporating the tax, she mines immediately if and only if it is socially efficient to do so. Note also that $V^*(0) = \hat{V}(0)$ because whether one learns or not, the decision to mine immediately returns the same expected payoff. Thus, we can represent the optimal mining tax explicitly as a function of the underlying parameters of the problem specified above, as is summarized by the following proposition:

Proposition 4. *Define a "informational tax" as*

$$Q = \hat{V}(\tau) - V^*(\tau) \tag{10}$$

Where:

$$V^*(\tau) = P(0) - c_w + e^{-r\tau}(c_w - \bar{c}_A)$$
$$\hat{V}(\tau) = \int_{-\infty}^{c_w} \left[P(0) - c_w + (c_w - c_A)e^{-r\tau}\right] f(c_A)dc_A + \int_{c_w}^{\tilde{p}} \left[P(0) - c_w + (c_w - c_A)e^{-rT}\right] f(c_A)dc_A$$

Such an informational tax, in addition to expected environmental costs, imposed on early extraction by a private mine owner, induces the economically efficient mine timing.

This core result defines an "information hurdle" that must be cleared in order to proceed with mining immediately. For instance, if a mine is currently viable based on private costs, but there are uncertain external costs to be learned in the future (e.g., environmental costs), then Equation 10 provides a way of determining the extra payoff necessary for a privately desirable mine to be socially desirable.

While useful for correcting behavior, the tax, Q, is only one possible tax from a family of taxes that would all convert a private mine owner into a socially efficient decision-maker. To see this, consider the simple case in which the rational planner is barely in favor of mining now rather than delay, with $\hat{V}(0) = \hat{V}(\tau) + \eta$, for small η. It follows that $V^*(0) > V^*(\tau)$.[9] In that case, even without a tax, the private owner will make the socially-desirable decision. Yet the tax, Q, can impose a large (and unnecessary) cost on the mine owner.[10] In this case, the private mine owner is made to pay a large mining tax even though the tax has no desired, or actual, behavioral effects.

To get around this problem, we derive an alternative tax, which we call the "minimum corrective tax," that is the smallest possible tax that just corrects the private mining behavior. The minimum corrective informational tax is given by:

$$Z = \begin{cases} V^*(0) - V^*(\tau) & \text{if} \quad \hat{V}(\tau) > \hat{V}(0) > V^*(\tau) \\ 0 \quad \text{otherwise} \end{cases} \tag{11}$$

In the example provided above (where $\hat{V}(0) = \hat{V}(\tau) + \eta$), no informational tax would be levied. In the numerical example that follows, we will calculate and compare the taxes Q and Z.

4.1 Illustrative example

We now illustrate the results of this analysis and provide a concrete example of the corrective mining taxes derived above. Let the initial reserves be given by R_0, and let market demand be a linear function of quantity extracted in a period: $P(t) = \alpha - \beta q(t)$. Global marginal extraction cost is c_w. We selected a set of parameters very loosely chosen to reflect the global market for gold.[11] Using these parameters, backward induction reveals a time to exhaustion of $T = 31$ years and a resulting initial price of $P(0) =$\$1,234. Now consider a deposit with an uncertain social cost of extraction (the sum of private and environmental costs), c_A. In particular, let the probability density function over true deposit social extraction cost c_A be given by $f(c_A) \sim N(\mu, \sigma^2)$, and assume that the true cost will be revealed in $\tau = 10$ years. Here, μ captures the expected social cost of mining. Because the naïve planner seeks to internalize expected social cost through a levy equal to expected environmental costs, the private mine owner faces cost, μ. In other words, the cost μ already captures *expected* environmental costs. We compare the private mine owner and rational planner decisions in this context parametrically, letting μ and σ vary.

The cost in the rest of the world is $c_w = 100$; if $\mu \geq 100$ (i.e. if expected cost in A is greater than or equal to the cost in the rest of the world), then both the mine owner and the sophisticated planner would find it efficient to postpone mining. But if $\mu < 100$, the private mine owner will mine immediately, even though the sophisticated planner might find it optimal to postpone mining until date $\tau = 10$ when the true environmental costs will be revealed. Panel (a) of Figure 1 illustrates how the decisions differ for the two types of decision-makers, over the parameter space for this problem (expected cost, μ, on the vertical axis, and uncertainty in that cost, σ, on the horizontal axis). In the top region (for $\mu > 100$) it is optimal for both the mine owner and sophisticated planner to delay mining. In the

[9]If a sophisticated decision maker is in favor of mining immediately, then a private mine owner is even more strongly in favor of mining immediately.

[10]The tax would be $Q = V^*(0) - V^*(\tau)$.

[11]The parameters are: $\alpha = 5e3$, $\beta = 5e3/2.5e8$, $c_w = 100$, $r = .05$, $R_0 = 3.58e9$.

lowest region, where expected cost and uncertainty are low, both decision makers would find it optimal to mine immediately. But in the middle region, the mine owner would like to mine immediately, while the sophisticated planner would delay mining until at least date τ. In this region, intervention will be required to incentivize the mine owner to delay mining in order to achieve efficiency.

Implementing either mining tax Q or Z induces the optimal behavior by the mine owner. The bottom panels of Figure 1 show the same parameter space as in the top panel, but overlay the two mining taxes (panel (b) shows the tax, Q, and panel (c) shows the minimal corrective tax, Z). Clearly Q is increasing in both μ and σ, and is larger than Z. Note that near the bottom of the figures (when expected cost is very low), it is optimal for both decision makers to mine immediately; in this region Q must be paid by the private mine owner, but doing so does not change her behavior (she still mines immediately). Near the top of the figures (when expected cost is very high), it is optimal for both decision makers to delay, so again Q does not change behavior. But for a band in the middle, Q causes the mine owner, who would have found it optimal to mine immediately, to delay mining until at least date τ. Notice that for any level of μ, the size of this region is increasing in σ (consistent with Corollary 1b).

We have already argued that Q, based on the quasi-option value, may be unnecessarily large to induce the behavior desired of the private mine owner. This is illustrated in panel (c) of Figure 1, where there are positive values of Q in all three regions of the figure, yet only in the middle wedge is corrective action needed – this is the only region in which there is a divergence between the private mine owner (for whom *expected* environmental costs have been priced in) and rational decision-making. But even if we only charged Q in the wedge between the vertical and slanted lines, it would still be excessive. This was the logic that underpinned our derivation of Z in Section 4. The tax Z is clearly non-zero only in the wedge, the only place where behavior needs correcting. For example, suppose the cost in region A is expected to be 20% lower than the cost in the rest of the world ($\mu = 80$ vs. 100), and suppose there is relatively high uncertainty about that cost ($\sigma = 100$). Under those parameters, the taxes are $Q =$\$12.30 and $Z =$\$7.60; since the initial price is \$1,234, these represent about a 1% tax, should the mine owner decide to mine immediately; this is sufficient to get the private mine owner to optimally delay mining.

Finally, Figure 2 provides an illustration of the effects of the pace of learning (τ) on behavior. Consistent with Corollary 1a, as learning is delayed (i.e. as τ increases), the parameter space over which intervention is required shrinks. Under these parameters, if learning will not take place for 30 or more years, no intervention is required to align the behavior of the two agencies. Put differently, the more proximally we expect to learn about the environmental costs of mining, the more likely we are to want to induce a delay by the private mine owner.

5 Extension: All Mines Face Uncertain Costs

So far we have focused on the case when a single, small mine owner faces uncertain environmental costs but where the rest of the competitive industry has certain costs. Here we put aside the issue of environmental costs and the distinction between the mine owner decision and the decision of the environmental planner. We ask what happens if *all* owners face uncertain private costs? To analyze this challenging problem we must make some concrete assumptions. First, we continue to assume that all learning occurs at date τ. Second, we assume that mining costs are uncertain and equal to ϕ, a random variable with mean c, and that no mine owner knows her true cost until date τ. This mirrors the model from above, but allows all owners to learn the true value of ϕ at date τ.

We must also specify the true underlying distribution of cost types. If we start by assuming that all mines face the same (albeit uncertain) cost, but that the cost is revealed only at date τ, we can immediately gain traction on the problem. In that case, we conclude the following:

Proposition 5. *If all mine owners in the market face the same, uncertain private cost, then the following results hold:*

(a) The price and extraction paths before τ are dependent only on c, the expectation of ϕ.

(b) The price path after τ depends on whether the true cost is revealed to be greater or less than c. If it turns out to be a higher cost, then the price discontinuously jumps up. If the it turns out to be a lower cost, then the price discontinuously jumps down.

(c) There is no incentive for delay to time τ.

Proposition 5 is useful in gaining some intuition about this extension, but does not portend very interesting behavior. Prior to τ, all owners are identical because they all face exactly the same beliefs about their costs. Interestingly, all owners also realize that after τ they will *also* be identical because they all have the same cost. While they do not know what that cost will be, there is no incentive to delay because revealing their true cost does not confer any benefit. Thus, this type of cost uncertainty (where all costs are uncertain, but identical) neutralizes the behavior identified in the rest of this paper.

Allowing all costs to be uncertain, but identical, is somewhat dissatisfying for a number of reasons. First, it seems unrealistic to assume that all mines receive the same cost shock. Rather, for example, mining in Alaska might be revealed to bring different costs than, say, mining in Peru. And second, if all costs were indeed the same, then why couldn't these costs be revealed upon the first instance of mining (by any firm) at the beginning of time? To overcome these objections, we now extend the model by assuming that each mine has cost that is drawn from a two-part i.i.d. distribution:

$$\phi = \begin{cases} c - \theta & \text{with probability} \quad 0.5 \\ c + \theta & \text{with probability} \quad 0.5 \end{cases} \tag{12}$$

In that setting, all owners face a strong incentive to delay until τ for the reasons underpinning this paper. If an owner delays until τ and is revealed to have high cost, then she can optimally delay until near final exhaustion of the resource, when her profits from mining will be maximized. If she delays until τ and is revealed to have low cost, then she can mine immediately. Thus, upon realizing her true cost, she can mine at the most economically efficient date. But if all owners delay, the pre-τ price would rise rapidly, and would presumably become sufficiently attractive to lure some owners to mine without the benefit of knowing their true costs. Taking all of this into account, we conclude the following:

Proposition 6. *If each mine owner faces an uncertain cost ϕ, drawn individually and independently from the same two-point distribution with mean c (and revealed at date τ), then:*

(a) The price path prior to τ is higher, and the quantity path prior to τ is lower, than under no uncertainty

(b) The price immediately and expectedly jumps downward at time τ.

(c) The price rises rapidly after τ until all low-cost reserves are exhausted. The price rises modestly thereafter until all high-cost reserves are exhausted.

The price and extraction behavior revealed by Proposition 6 is strikingly different than what the standard Hotelling model predicts. Proposition 6(a) is consistent with our intuition because mine owners would like to delay, which drives extraction down and prices up. But Proposition 6(b) reveals that all owners expect the price to immediately and discontinuously jump down at date τ. Why, then, wouldn't an owner who mines just after τ move her extraction slightly backward in time to take advantage of the high price? The answer is that she would, if she knew she had low cost. But because no owner knows her cost in advance, the intertemporal arbitrage condition states that the *expected* present value of rent must be identical whether one commits to mining prior to τ, or waits until after τ and then sequences herself to maximize profit.

The equilibrium price path that is the outcome of this option value behavior is depicted in Figure 3, assuming a simple high-low cost distribution where $c = 100$, $\theta = 2000$, and all parameters conform to the previous numerical illustration. Here, the pre-τ price is substantially higher with uncertainty (red curve) than without uncertainty (black curve). The price immediately jumps down at time $\tau = 10$. The remaining low-cost reserves are extracted between periods 10-17.5, during which time prices rise rapidly because cost is so low. After time 17.5, only high-cost ore is available. The price path shifts, but cannot jump because this would introduce an arbitrage opportunity. In this simulation, uncertainty extends the time over which mining occurs from 30 years (in the no-uncertainty case) to 36.5 years.

6 Discussion

The main result of this paper is that *any* uncertainty about the environmental cost of mining can lead to an efficient social decision to postpone extraction for a marginal mine, even when the *expected* environmental costs suggest that it is optimal to mine immediately. However, this doesn't imply the mine will *never* be exploited. Rather, as information on environmental cost is revealed, it gives rise to new socially optimal decisions about mine-timing. We showed that if the cost is learned to be relatively low (or even as "expected") then it will be optimal to mine once that information is revealed. On the other hand, if cost is revealed to be large (but not too large), then it will be optimal to mine at a later date. In extreme cases in which the revealed cost turns out to be very large, it will be optimal to postpone indefinitely.

Our focal application is when all uncertainty regards external costs, such as environmental costs. Internalization of those costs is manifest in a regulatory authority granting or denying permission to mine. If the regulatory authority does a standard cost-benefit analysis, based on expected external costs, then it is acting naïvely, and may erroneously grant the right to mine immediately when it is optimal to delay mining (the red wedge in Figure 1(a)). In such a setting, we showed that an additional cost, equal to the mining tax, must be imposed on private mine owners to induce efficient mining behavior.

Our other intended contribution is that we provide the first theoretical treatment of the equilibrium outcome of a decentralized economy in which all agents face an incentive to delay. In that setting, some owners do indeed delay, which drives up pre-learning prices. But we found that prices expectedly and abruptly jump down immediately upon learning. That part of the model reveals price dynamics that are fundamentally different that are expected from both the standard Hotelling model, and the model of mining under uncertainty, but without learning. Thus, the possibility of learning can qualitatively alter the equilibrium price paths we observe in natural resource markets, and may help resolve long-standing debates about why natural resource prices fail to reflect what standard Hotelling theory predicts.

Incorporating uncertainty and learning into a Hotelling model reveals that delay of extraction occurs more often than expected. This is the case with uncertainty for a single mine as well as uncertainty over costs for multiple mines. The implication of deferring extraction is that the price path over time tends to flatten – current prices increase and future prices decrease. This suggests that quasi-option value may be one reason exhaustible resource prices do not rise as quickly as suggested by the simple Hotelling model.

One extension of this model, which we do not analyze here, occurs when the regulatory authority is subject to political pressure. In that case it may be in the mine's interest to lobby the regulatory authority to grant permission to mine prior to uncertainty being resolved. The framework introduced in this paper would allow us to compute the private mine owner's willingness to pay to obtain early approval rather than wait until uncertainty is resolved and run the risk that external costs end up being high, resulting in a denial of permission to mine.

Our simple theoretical treatment sheds light on a rich array of empirically and policy relevant contemporary problems involving the extraction of non-renewable resources with uncertainty about the associated environmental costs. In such real world applications, environmental costs are often pitted against extraction benefits. This analysis reveals that this is a false tradeoff: The question is not whether we should *ever* mine, but rather whether we should mine today or postpone the decision until better information on environmental costs is revealed. After all, a decision to delay mining forfeits none of the natural resource - it just saves it for the future when a more insightful decision can be made.

We now briefly return to our motivating example of the Alaskan Pebble Mine. The analysis in this paper begs the question: Have the planners in charge of permitting the mine accounted for the exhaustible nature of the resources, the price paths of those resources and the nature of uncertainty and learning? This paper provides a framework to rationally structure our thinking about such challenges.

Appendix of Proofs

Proof to Proposition 1

Proof. Let V_1 be the expected payoff from mining prior to τ and let V_2 be the expected value of waiting until τ, learning the true value of c_A, and then deciding when to mine. To obtain V_1, because $\bar{c}_A = c_w$, Equation 7 reveals that the non-learning mine owner (who faces $\bar{c}_A$) is indifferent between mining at any date $t \in [0, T]$; in other words, she obtains the same expected payoff from extracting at any date in that closed interval. Since the mining date is irrelevant, we will use date τ for convenience. Using Equation 2,

$$V_1 = \int_{-\infty}^{\infty} \left[P(0) - c_w + (c_w - c_A)e^{-r\tau} \right] f(c_A) dc_A \tag{13}$$

which we split into the relevant regions of the realization of c_A:

$$\begin{aligned} V_1 = & \int_{-\infty}^{c_w} \left[P(0) - c_w + (c_w - c_A)e^{-r\tau} \right] f(c_A) dc_A \\ & + \int_{c_w}^{\tilde{p}} \left[P(0) - c_w + (c_w - c_A)e^{-r\tau} \right] f(c_A) dc_A \\ & + \int_{\tilde{p}}^{\infty} \left[P(0) - c_w + (c_w - c_A)e^{-r\tau} \right] f(c_A) dc_A \end{aligned} \tag{14}$$

V_2 can be decomposed in a similar manner, though the timing of mining will depend on the realized value of c_A. We have

$$\begin{aligned} V_2 = & \int_{-\infty}^{c_w} \left[P(0) - c_w + (c_w - c_A)e^{-r\tau} \right] f(c_A) dc_A \\ & + \int_{c_w}^{\tilde{p}} \left[P(0) - c_w + (c_w - c_A)e^{-rT} \right] f(c_A) dc_A \\ & + 0 \end{aligned} \tag{15}$$

where the third term is 0 because no mining takes place in the event that $c_A > \tilde{p}$.

We would like to prove that $V_2 > V_1$. Taking the difference, we see:

$$V_2 - V_1 = \underbrace{\int_{c_w}^{\tilde{p}} (c_w - c_A)(e^{-rT} - e^{-r\tau}) f(c_A) dc_A}_{Term1} - \underbrace{\int_{\tilde{p}}^{\infty} \left[P(0) - c_w + (c_w - c_A)e^{-r\tau} \right] f(c_A) dc_A}_{Term2} \tag{16}$$

The first term on the RHS is unambiguously positive (this is the benefit of mining at T rather than τ when the true cost is in $[c_w, \tilde{p}]$). The second term is unambiguously negative (it is the cost of mining at τ when the cost is extremely high ($c_A > \tilde{p}$). Thus, $V_2 > V_1$, which concludes the proof. □

Proof to Proposition 2

Proof. Let V_1 be the expected payoff of committing to mine at some date $s < \tau$ and let V_2 be the expected payoff from postponing the decision until at least date τ. Finally, let V_T

be the expected value of committing to mine at date T. By Equation 7, $V_T > V_1$ (when $c_w < \bar{c}_A < \tilde{p}$ the non-learning mine operator (who faces $\bar{c}_A$) maximizes payoff by mining at date T). We will show that $V_2 > V_T$, which implies that $V_2 > V_1$. We have:

$$\begin{aligned} V_1 < V_T = & \int_{-\infty}^{c_w} \left[P(0) - c_w + (c_w - c_A)e^{-rT}\right] f(c_A)dc_A \\ & + \int_{c_w}^{\tilde{p}} \left[P(0) - c_w + (c_w - c_A)e^{-rT}\right] f(c_A)dc_A \\ & + \int_{\tilde{p}}^{\infty} \left[P(0) - c_w + (c_w - c_A)e^{-rT}\right] f(c_A)dc_A \end{aligned} \tag{17}$$

And V_2 is given in Equation 15. The difference is given by:

$$V_2 - V_T = \underbrace{\int_{-\infty}^{c_w} (c_w - c_A)(e^{-r\tau} - e^{-rT}) f(c_A)dc_A}_{\text{Term 1}} - \underbrace{\int_{\tilde{p}}^{\infty} \left[P(0) - c_w + (c_w - c_A)e^{-rT}\right] f(c_A)dc_A}_{\text{Term 2}} \tag{18}$$

Term 1 is clearly positive and Term 2 is clearly negative, so $V_2 > V_T > V_1$, and $V_2 > V_1$ which proves the result. □

Proof to Proposition 3

Proof. If $\bar{c}_A < c_w$ and A commits to mining prior to τ, she should mine at date 0 (see Equation 7). Let V_1 be the expected payoff from doing so. Let V_2 be the expected payoff from delaying the decision until date τ. We have:

$$V_1 = \int_{-\infty}^{c_w} [P(0) - c_A] f(c_A)dc_A + \int_{c_w}^{\tilde{p}} [P(0) - c_A] f(c_A)dc_A + \int_{\tilde{p}}^{\infty} [P(0) - c_A] f(c_A)dc_A \tag{19}$$

And V_2 is given in Equation 15. The difference is given by:

$$\begin{aligned} V_2 - V_1 = & \underbrace{\int_{-\infty}^{c_w} (c_w - c_A)(e^{-r\tau} - 1) f(c_A)dc_A}_{\text{Term 1}} \\ & + \underbrace{\int_{c_w}^{\tilde{p}} (c_w - c_A)(e^{-rT} - 1) f(c_A)dc_A}_{\text{Term 2}} \\ & + \underbrace{\int_{\tilde{p}}^{\infty} [c_A - P(0)] f(c_A)dc_A}_{\text{Term 3}} \end{aligned} \tag{20}$$

Term 1 is negative, and Terms 2 and 3 are positive. Proving that conditions exist under which it is optimal to postpone the decision until τ requires showing conditions under which $V_2 > V_1$; we use the following sufficient condition. Take the limit of 20 as $\tau \to 0$. Doing so only affects Term 1, and has no effect on Terms 2 or 3. Clearly $\lim_{\tau \to 0}$ Term 1 $= 0$, while Terms 2 and 3 are strictly positive. Thus, small τ provides a sufficient condition for the result to hold. See also the numerical example for a much larger parameter space over which this result holds. □

Proof to Corollary 1

Proof. **Corollary 1(a)**: Let $\Delta \equiv V_2 - V_1$ from Equation 20. Taking the derivative gives:

$$\frac{d\Delta}{d\tau} = \int_{-\infty}^{c_w} -r(c_w - c_A)e^{-r\tau} f(c_A) dc_A < 0 \tag{21}$$

which concludes the proof.
Corollary 1(b): Define $\Phi(c_A)$ as the difference in payoffs between postponing the decision until τ and mining prior to τ, when the true value of A's cost parameter is c_A. First note that $\Delta \equiv V_2 - V_1$ is the integral of $\Phi(c_A)$, weighted by the probability density, $f(c_A)$, as follows.

$$\Delta = V_2 - V_1 = \int_{-\infty}^{\infty} \Phi(c_A) f(c_A) dc_A \tag{22}$$

To prove that Δ is increasing in the "risk" of $f(c_A)$ (colloquially, in the "spread" of $f(c_A)$), we rely on the main result of Rothschild and Stiglitz (1970), implying it is sufficient to show that $\Phi(c_A)$ is a convex function of c_A. The function $\Phi(c_A)$ is given as follows:

$$\Phi(c_A) = \begin{cases} (c_w - c_A)(e^{-r\tau} - 1) & \text{if } c_A < c_w \\ (c_w - c_A)(e^{-rT} - 1) & \text{if } c_w \leq c_A \leq \tilde{p} \\ c_A - P(0) & \text{if } c_A > \tilde{p} \end{cases} \tag{23}$$

Thus, $\Phi(c_A)$ is a continuous, piecewise linear, increasing function of c_A, where each subsequent line segment has a higher slope, thus, $\Phi(c_A)$ is convex and the result is proven. $\square$

Proof to Proposition 5

Proof. Prior to τ, only the expected value of cost is known to each mine owner. There is no strategic advantage to be gained by slowing or accelerating depletion to take advantage of full information at τ, since all producers have the same cost. With no advantage to deferring to τ, rents based on expected costs will rise at the rate of interest until τ, at which point prices will either drop or rise, depending on the revealed cost, and from that point on rents will continue to rise at the rate of interest. An additional constraint on the price path is that the expected price just after τ (based on information just prior to τ) must be the same as the actual price just prior to τ. Otherwise there would be arbitrage opportunities to restore balance. $\square$

Proof to Proposition 6

Proof. During the period prior to τ, rents must rise at the rate of interest. During this period, no producer would have an incentive to defer to τ, since all producers would have that same incentive. At τ, when costs are fully revealed, producers sort themselves, with lowest cost producers (those with the highest rent) mining first, following sequentially by higher and higher cost producers. Those conditions, along with resource quantity constraints, uniquely defines the price path articulated in the proposition. $\square$

References

Arrow, K. J. (1968). *Optimal Capital Policy and Irreversible Investment.*

Arrow, K. J. and A. C. Fisher (1974). Environmental preservation, uncertainty, and irreversibility. *The Quarterly Journal of Economics*, 312–319.

Arrow, K. J. and M. Kurz (1970). Optimal growth with irreversible investment in a Ramsey model. *Econometrica 38*, 331–44.

Dennis, B. (July, 2017). Controversial Alaskan gold mine could be revived under Trump's EPA. *The Washington Post*.

Dixit, A. K. and R. S. Pindyck (1994). *Investment under Uncertainty.* Princeton University Press.

Environmental Protection Agency (2014). Proposed determination of the U.S. EPA Region 10 pursuant to section 404(c) of the Clean Water Act, pebble deposit area, southwest alaska. Technical report, US EPA.

Grandoni, D. and J. Partlow (2021). Epa to protect Alaska's Bristol Bay, blocking major gold mine. *The Washington Post*.

Hanemann, W. M. (1989). Information and the concept of option value. *Journal of Environmental Economics and Management 16*(1), 23–37.

Henry, C. (1974). Option values in the economics of irreplaceable assets. *The Review of Economic Studies*, 89–104.

Hoel, M. (1978). Resource extraction, uncertainty, and learning. *The Bell Journal of Economics*, 642–645.

Holley, E. A. and C. Mitcham (2016). The pebble mine dialogue: A case study in public engagement and the social license to operate. *Resources Policy 47*, 18–27.

Hotelling, H. (1931). The economics of exhaustible resources. *Journal of political Economy 39*(2), 137–175.

Mensink, P. and T. Requate (2005). The Dixit-Pindyck and the Arrow-Fisher-Hanemann-Henry option values are not equivalent: a note of Fisher (2000). *Resource and Energy Economics 27*, 83–8.

Narula, S. K. (2014, March 14). Is Alaska's Pebble Mine the Next Keystone XL? *The Atlantic*.

Parker, G. Y., F. M. Raskin, C. A. Woody, and L. Trasky (2008). Pebble mine: Fish, minerals, and testing the limits of Alaska's large mine permitting process. *Alaska L. Rev. 25*, 1.

Pindyck, R. S. (1980). Uncertainty and exhaustible resource markets. *The Journal of Political Economy*, 1203–1225.

Rothschild, M. and J. E. Stiglitz (1970). Increasing risk: I. a definition. *Journal of Economic theory 2*(3), 225–243.

Traeger, C. P. (2014). On option values in environmental and resource economics. *Resource and Energy Ecponomics 37*, 242–52.

Weisbrod, B. (1964). Collective-consumption services of individual-consumption goods. *Quarterly Journal of Economics 78*, 471–77.

Figure 1: Each panel shows the parameter space over which different outcomes arise. Panel (a) shows the decisions of mine Owner and Sophisticated planner. Panels (b) and (c) show corrective taxes, Q and Z.

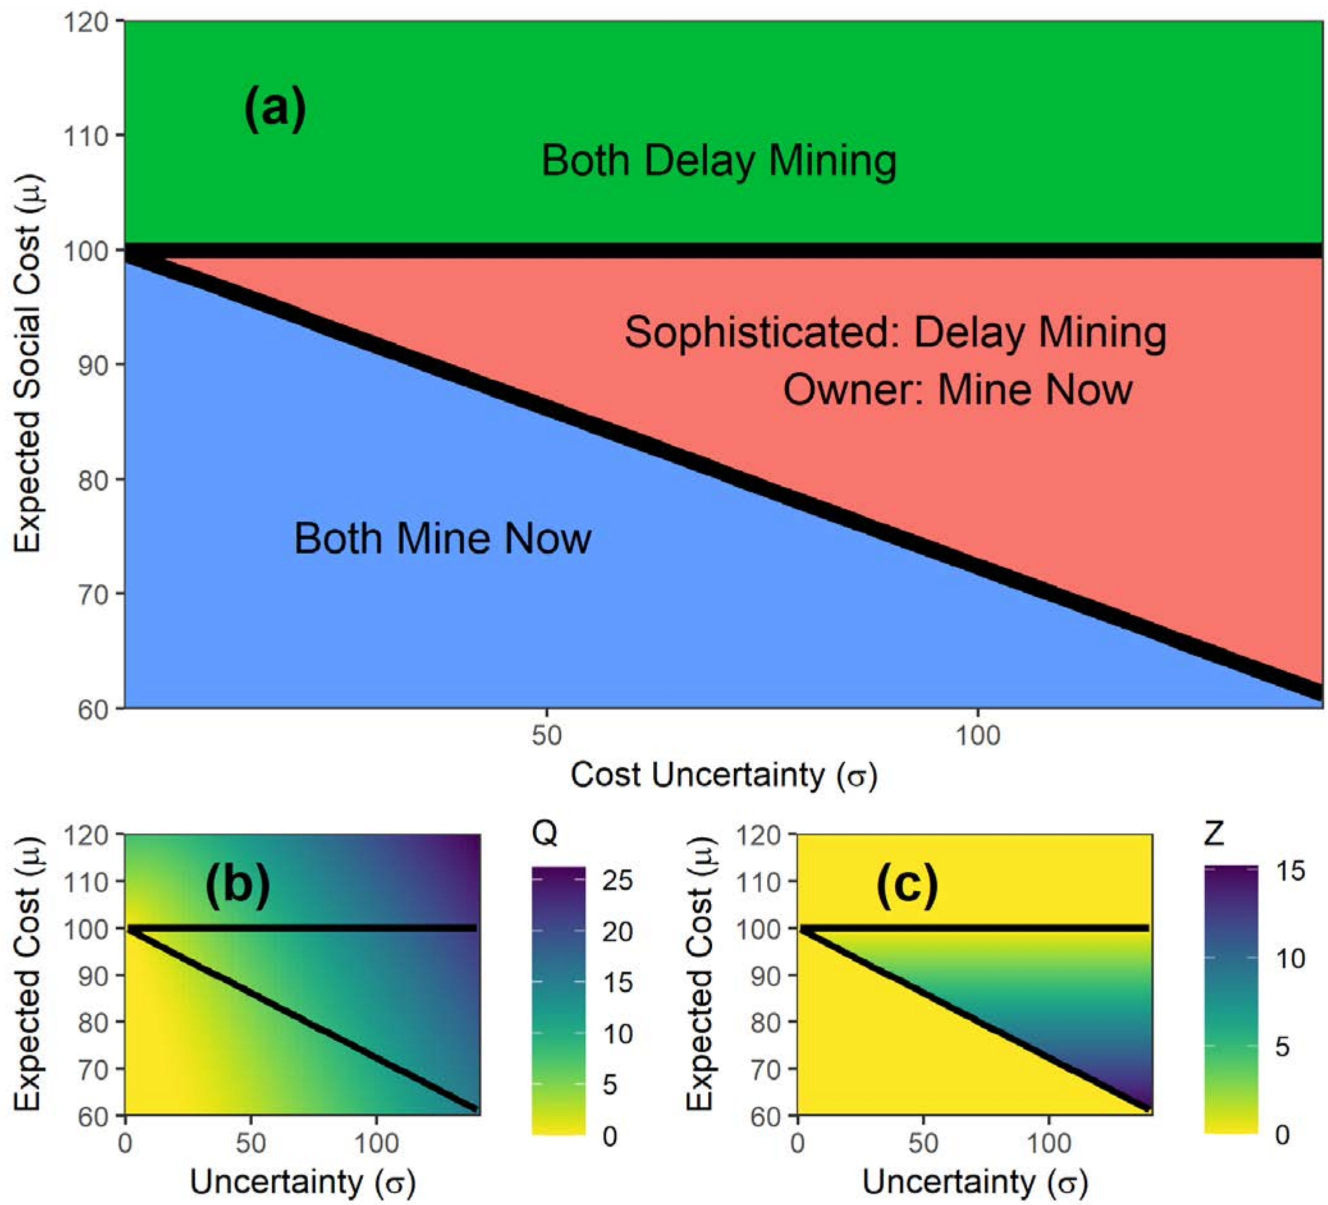

Figure 2: Panels show parameter space for increasing values of τ. Consistent with Corollary 1, the space over which intervention is required is shrinking in τ.

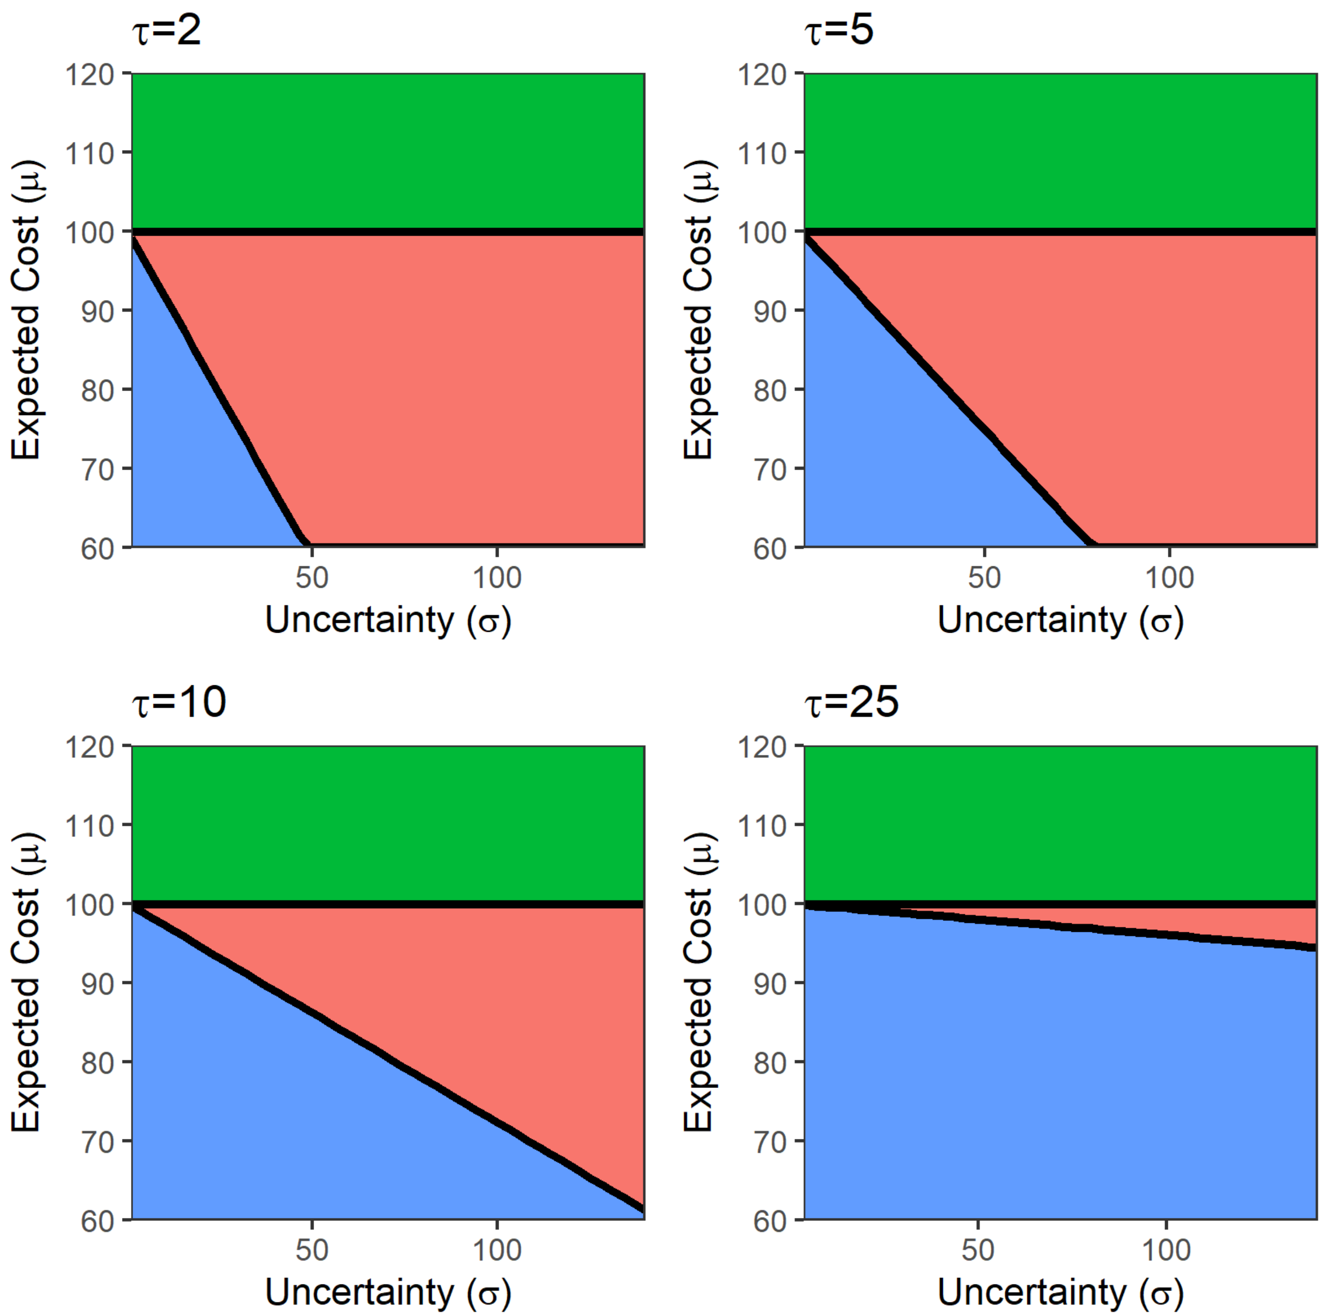

Figure 3: Equilibrium price path under i.i.d. uncertainty, as described in Equation 12. The shaded path is with uncertainty. The solid black path is without uncertainty. Information is resolved at date $\tau = 10$ (dashed vertical line). Low-cost ore is exhausted at time 17.5 (dashed vertical line).