Ecological Economics 239 (2026) 108750

Contents lists available at ScienceDirect

Ecological Economics

journal homepage: www.elsevier.com/locate/ecolecon

ANALYSIS

# The welfare gains from diversified environmental policies

Ben Blachly [a], Charles Sims [b,*], Travis Warziniack [c]

[a] Industrial Economics, Inc., Cambridge, MA 02140, United States
[b] Department of Economics, University of Tennessee, Knoxville, TN 37996, United States
[c] USDA Forest Service, Rocky Mountain Research Station, Fort Collins, CO 80526, United States

ARTICLE INFO

*JEL codes:*
Q57
Q51
D81

*Keywords:*
Choice experiment
Endangered species
Mean-variance utility
Modern portfolio theory
Risk aversion

ABSTRACT

This paper provides a utility-theoretic framework for allocating a budget across multiple environmental projects with uncertain outcomes. We utilize portfolio theory to derive the price of a risk (i.e., variance) reduction, then we estimate willingness to pay for the same risk reduction by analyzing stated preference data using a mean-variance specification for utility. The welfare-maximizing budget allocation is determined by equating price and willingness to pay. We use our framework to make two points. First, the welfare gains from a diversified environmental policy (i.e., one composed of multiple environmental projects) depend on society's willingness to pay for policy risk reduction and the price of a risk reduction determined by the covariance between projects. In an application to salmon restoration in Maine, we find that the welfare gains from policy diversification are equal to 5 % of society's willingness to pay for salmon restoration overall. In addition, failing to account for aversion to outcome variance may lead to underestimates of willingness to pay for environmental projects and over-investment in the riskiest projects. Failing to account for aversion to variance in the number of salmon underestimates the welfare associated with salmon restoration by over 80 % and funnels all restoration into the watershed where successful restoration is most uncertain.

The adoption of environmental policies imposes costs on society in exchange for a social return in the form of a stream of benefits such as improved health outcomes or reductions in a pollutant. However, these benefits are highly uncertain due to poor understanding of the biophysical processes connecting environmental policies to outcomes and the inherent stochasticity of natural processes (Pindyck, 2007). Further, policy implementation implies choosing among several candidate projects each with a different level of risk and return. Examples include allocating pollution control across multiple point sources, catch limits for multiple fisheries, and habitat restoration across multiple sites. Risk aversion implies that the project that generates the highest expected return may not always be preferable if it is also relatively risky. Further, if the benefits of projects covary, risk can be reduced (i.e., diversified) with a policy portfolio whereby multiple projects receive investment instead of investing in a single project. This paper develops a novel utility-theoretic framework for allocating investments across multiple uncertain environmental projects and uses that framework to provide the first estimates of welfare gains from risk diversification in environmental policy.

Our framework extends and synthesizes two veins of the environmental economic literature: the growing body of portfolio theory applications to environmental management and the non-market valuation literature on valuing environmental projects with uncertain outcomes. Portfolio theory formalizes the notion that overall risk—measured by variance—can be reduced by holding multiple uncorrelated assets in a portfolio (Markowitz, 1952). In the context of environmental management, it has been applied to inform forest management (Crowe and Parker, 2008), fisheries (), wetland conservation (Ando and Mallory, 2012; Shah and Ando, 2015; Vinent et al., 2019), and marine spatial planning (Halpern et al., 2011) and water management (Leroux and Martin, 2016; Leroux et al., 2018; von Eiff and Pommeret, 2021; von Eiff et al., 2023; Leroux et al., 2022).[1] Most of these applications assume the asset composition of the portfolio is static. However, a handful of applications (e.g., Leroux et al., 2022) identify an optimal dynamic portfolio that jointly optimizes the consumption path from a portfolio of assets with stochastic returns and the asset composition (weights) of the portfolio over time (Merton, 1969, 1971).

Ando et al. (2018) evaluated 26 environmental MPT applications to

* Corresponding author.
*E-mail address:* cbsims@utk.edu (C. Sims).

[1] See Matthies et al. (2019) for a review.

https://doi.org/10.1016/j.ecolecon.2025.108750
Received 27 January 2025; Received in revised form 27 July 2025; Accepted 29 July 2025
Available online 6 August 2025
0921-8009/

explore the circumstances under which diversification was an effective strategy to mitigate risk. Using curvature of the efficient frontier as the primary metric for diversification potential, they showed that diversification is most effective in three cases: (1) when multiple assets are negatively correlated across the risk domain; (2) when the two highest return assets have similar expected values; and (3) when multiple assets carry low risk. While this work identifies the price of risk diversification (i.e., the foregone return), it does not estimate the benefits, measured by society's willingness to pay for risk diversification. To determine the optimal level of diversification, previous environmental applications have either utilized a simple criteria like maximizing the return/risk ratio (Halpern et al., 2011), assumed a risk aversion parameter (Vinent et al., 2019), or calibrated a risk aversion parameter from data on resource management decisions (Leroux and Martin, 2016; Leroux et al., 2018). We extend this literature by showing how commonly used choice experiments can be used to estimate the degree of risk aversion, identify a welfare-maximizing portfolio, and estimate the welfare gains from risk diversification.

The non-market valuation literature has shown that treating risky outcomes as certain within a valuation task can lead to can lead to empirical value estimates that do not reflect ex ante welfare change under uncertainty since individuals often require a risk premium when purchasing risky assets (Johnston et al., 2017).[2] A growing number of stated preference studies, therefore, present respondents with a probability of policy success (or failure) (e.g., Roberts et al., 2008; Glenk and Colombo, 2013; Bartczak et al., 2015; Rolfe and Windle, 2015; Makriyannis et al., 2018; Faccioli et al., 2019; Meles et al., 2023). The social value of these uncertain environmental policies (measured by society's willingness to pay for them) is driven by society's preferences for risk and reward which govern tradeoffs between the policy's expected benefits and the variance over those benefits (Pindyck, 2014). Of the empirical models appearing in the literature, however, most are linear in attributes and do not include a measure of the outcome variance, implicitly ignoring the risk-return tradeoffs implied by risk aversion.[3]

An exception is Holzer and McConnell (2017) who specify and estimate different models of constant absolute risk aversion (CARA) to estimate welfare associated with recreational fishing trips. We build on Holzer and McConnell's work by showing how CARA utility functions generate risk preferences consistent with mean-variance utility (i.e., individuals have preferences for the mean and variance of policy outcomes). We then use our mean-variance utility framework to show that failing to account for an individual's aversion to variance in the policy outcomes will often lead to underestimates of the willingness to pay for the project and over-investment in the riskiest environmental projects. One advantage of our mean-variance approach is that it allows for direct estimation of risk preferences from commonly used choice experiments.

To illustrate, we use data from a choice experiment designed to elicit preferences for restoring Atlantic salmon habitat in the state of Maine, U. S. Atlantic salmon have been an endangered species since 2000. Ongoing and future recovery efforts are expensive (NOAA, 2019) and subject to a high degree of outcome uncertainty (Mills et al., 2013). We present respondents with choices between hypothetical restoration scenarios consisting of an estimate for salmon carrying capacity alongside a probability that salmon will utilize the new habitat. Consistent with much of the existing empirical literature, we first model the data by specifying utility over the expected number of salmon only. We then include the implied variance. A likelihood ratio test indicates that the model that includes variance fits the data significantly better. We find an average WTP of $21.50 to reduce outcome variance by 1000 salmon, holding the expected number of salmon constant. The commonly-used choice model that ignores variance underestimates the welfare associated with restoration in each of the Maine watersheds we consider by over 80 %. This supports the theoretical conclusion of Pindyck (2014), which shows that risk aversion implies bias associated with policy analyses based on expected return only.

Next, we estimate the expected number of salmon and the associated risk (variance) resulting from habitat restoration in six Maine watersheds by pairing a habitat model with historic annual salmon counts. Considering each watershed as a potential restoration asset, we utilize MPT to identify the restoration portfolios (i.e., budget allocations across the six watersheds) that maximize return for each level of risk. These efficient portfolios form a frontier, the slope of which indicates the 'price' of reducing risk. We use results from the choice experiment to estimate mean-variance indifference curves, which represent willingness to pay for a risk (variance) reduction. We then identify the welfare-maximizing restoration portfolio (i.e., the amount of the budget allocated to each watershed) by the tangency condition where marginal willingness to pay equals the price of reducing risk. We find a welfare improvement associated with diversifying restoration across two watersheds (Penobscot and Maine Coastal) compared to concentrating solely on the watershed with the highest expected return (Penobscot). Specifically, about 5 % of the welfare from salmon restoration ($25.04) can be attributed to the value society holds for reducing the riskiness of salmon restoration. We also find that a common risk diversification rule-of-thumb (i.e., evenly divide the restoration budget across all watersheds) reduces welfare by nearly 40 %.

## 1. Policy portfolio based on risk, return, and willingness to pay

A manager is given a fixed budget to implement a policy. A policy is an allocation of a fixed budget across multiple projects with uncertain returns. Projects might be, for example, various pollution controls at a point source or possible sites for wildlife habitat restoration, and a policy might be an allocation of pollution control expenditures across point sources or an allocation of habitat restoration funds across restoration sites. Let $N$ denote the manager's choice set of projects. The expected return of project $i \in N$ is $\mu_i$, and the variance of the return of that project is $\sigma_i{}^2$. The variance reflects the degree of uncertainty in project returns. This uncertainty could arise from unexpected changes in economic, political, and environmental factors that influence project returns over the long term.

Thus, policies are portfolios of projects that a manager chooses to invest in. The expected return and risk (variance) of this policy portfolio depends on how the budget is allocated, defined as a vector of weights, $w_i$ that represent the proportion of the fixed budget allocated to project $i$. A project is included in the policy if $0 < w_i \leq 1$ and excluded from the policy when $w_i = 0$. Our approach is fundamentally static in that the weights do not change over time. If economic, political, or environmental dynamics are critical for understanding project returns (and if the data is available to confidently specify these dynamic processes), a dynamic MPT approach can be applied to produce time-varying portfolio weights (e.g., Leroux et al., 2022).

Expected policy return is the weighted sum of the individual project returns: $\mu = \sum_i w_i \mu_i$. Policy variance is the weighted sum of individual project standard deviations and covariances: $\sigma^2 = \sum_i (w_i \sigma_i)^2 + \sum_i \sum_j (w_i w_j COV_{ij}) = \sum_i (w_i \sigma_i)^2 + \sum_i \sum_j (w_i w_j \sigma_i \sigma_j \rho_{ij})$, where $COV_{ij}$ is the covariance of projects $i$ and $j$, $\sigma_i$ is the standard deviation of the return from project $i$ and $\rho_{ij}$ is the correlation between projects $i$ and $j$. Hence if the returns are less than perfectly correlated ($\rho_{ij} < 1$), a portfolio of the projects will have a lower variance than either of the projects individually.

The covariance between projects implies a tradeoff between risk and return. A manager could choose to allocate the entire budget to the

[2] There are several other reasons risk may lead to biased welfare estimates including a respondent's subjective beliefs, an anchoring effect when respondents are unsure of the true value they place on environmental amenities, and ineffective communication of risks in stated preference questionnaires.

[3] de Palma et al. (2008) provides a review of discrete choice models under uncertainty.

project that provides the largest expected return. Alternatively, the manager could choose to allocate part of the budget to projects that are not perfectly correlated or are negatively correlated with the project with the highest expected return. A policy portfolio composed of several projects will have a lower return but will also be less risky than a policy portfolio composed of the single high-return project. A well-known result from the portfolio theory literature is that this risk-return trade-off can be represented by a concave frontier in $\mu_i$ - $\sigma_i{}^2$ space (Merton, 1971).

The MPT approach identifies this frontier by finding weights $\boldsymbol{w}$ that maximize expected returns, $\mu$, for a given level of risk or minimize risk for a given level of expected returns:

$$\min \boldsymbol{w}^T \Omega \boldsymbol{w} \tag{1}$$

subject to the policy budget being fully allocated ($\sum_i w_i = 1$) and $E[\mu]^T \boldsymbol{w} = \overline{\mu}$ where $\boldsymbol{w}$ is a vector with elements $w_i$, $\Omega$ is the covariance matrix of project returns, $\overline{\mu}$ is an exogenously determined target return for the policy, and $E$ and $T$ are expectation and transpose operators respectively. The location and curvature of the policy frontier can be found by solving the problem in Eq. (1) for each $\overline{\mu}$ within a discretized interval. We represent this policy frontier as $\mu = Y(\sigma^2)$ where $Y' > 0$ and $Y'' < 0$. Fig. 1 shows a hypothetical frontier. Policies that are to the right of this frontier are inefficient in that an alternative budget allocation could provide a higher expected return for the same level of risk (or the same expected return at a lower level of risk). Policies to the left of this frontier are not possible given the risk and return of the individual projects. The policy furthest right on the frontier allocates the entire budget to the project with the highest expected return. All other portfolios on the frontier are diversified in that they achieve a lower level of risk by allocating a portion of the budget to projects that are not perfectly correlated or negatively correlated.

### 1.1. Welfare-maximizing portfolio choice

While MPT can identify the frontier of preferred policy portfolios, it cannot identify the welfare-maximizing policy on the frontier. In financial applications of MPT, the optimal portfolio is often identified as the portfolio where the Capital Market Line (CML) is tangent to the efficient frontier. This approach makes it possible to identify a single optimal portfolio along the frontier without knowing the underlying utility function and risk preferences (Sharpe, 1964). We do not adopt this CML approach for two reasons. First, the CML implies the manager is willing and able to shift investment from environmental projects to risk- free assets, such as a U.S. Treasury bond or investments in well-diversified portfolios of assets other than environmental projects. This is typically not the case for environmental managers who are obligated to spend their budget exclusively on environmental projects. Second, the CML approach does not identify welfare-maximizing portfolios and thus cannot be used to identify the welfare benefits of diversification.

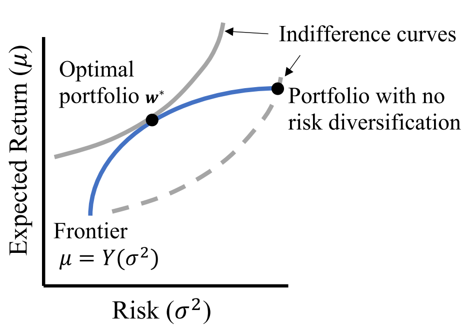


**Fig. 1.** Optimal policy portfolio $w^*$ allocates the budget to equate society's marginal rate of substitution between risk and return and the marginal rate of transformation between risk and return.

To identify the welfare-maximizing policy portfolio, we must assume that the manager has preferences for the mean and variance of the return generated from the policy portfolio. There are many reasons to believe that decision makers may be averse to variance. For example, more variable wildlife populations may lead to population collapse. Temporary high pollution concentrations can be as important for determining health impacts as long-run averages. Furthermore, government agency officials may be concerned that variability in wildlife populations or pollution levels may anger constituents and potentially lead to loss of their job.

We use a mean-variance utility function (Tobin, 1958) that gives the manager's random future real wealth $W = \mu + \sigma Z$, where $Z$ is a standardized random variable with density $f_Z(z)$. The manager has preferences over the risk and expected return of the policy portfolio. If the manager is risk averse, these preferences can be represented by upward-sloping indifference curves in $\mu_i$ - $\sigma_i{}^2$ space. To capture these risk preferences, the manager's expected utility is expressed as an indirect utility function

$$EU(W) = V(\mu, \sigma) = \int_{-\infty}^{\infty} U(\mu + \sigma z) f_Z(z) dz \tag{2}$$

with indifference curves that are upward sloping (Tobin, 1958).[4] Preferences for the mean and variance of the return generated from the policy portfolio is consistent with expected utility theory in three cases: 1) when the return distribution is normal, 2) when the underlying utility function is quadratic, and 3) for small risks.

When the distribution of returns is normal, expected utility is only a function of mean and variance. Consider, for example, an absolute risk aversion utility function $u(x) = 1 - \exp(-\alpha x)$, where $x$ is the return from the policy portfolio and $\alpha > 0$ is the coefficient of absolute risk aversion. If $x \sim N(\mu, \sigma^2)$, expected utility is $E[u] = 1 - \exp\left[-\alpha\mu + \frac{\alpha^2\sigma^2}{2}\right]$ since the certainty equivalent is $\mu - \alpha\sigma^2/2$.

When the underlying utility function is a quadratic function of the return, such as $u(x) = x - bx^2$ with $b < 0$, taking the expected value yields $E[u] = E[x] - bE[x^2]$.[5] The variance of random variable $x$ is $\sigma^2 = E[x - E(x)]^2$, which can be rearranged to show that $E[x^2] = \sigma^2 + (E[x])^2$. Substituting that expression for $E[x^2]$ in the expected utility function, and rewriting $E[x]$ as $\mu$ yields:

$$E[u] = E[x] - bE[x^2] = \mu - b(\sigma^2 + \mu^2) \tag{3}$$

Thus, quadratic utility implies that expected utility is solely a function of the mean and variance of $x$ regardless of the distribution of $x$.

When risks are small, the second order (local) Taylor series approximation of the von Neuman Morgenstern (VNM) utility function around the expected value is

$$u(x) = u(\mu) + u'(\mu)(x - \mu) + u''(\mu)\frac{(x - \mu)^2}{2} \tag{4}$$

with expected utility

$$E[u(x)] = u(\mu) + u''(\mu)\frac{E(x - \mu)^2}{2} = u(\mu) + u''(\mu)\frac{\sigma^2}{2}. \tag{5}$$

[4] The convexity of the indifference curves holds only if risk is measured in terms of standard deviation. In the ($\sigma^2$,$\mu$)-plane, the indifference curves are still upward sloping but may have either curvature.

[5] Many objections to the use of quadratic utility such as the implicit assumption of increasing absolute risk aversion are predicated on assuming a quadratic function of wealth instead of return. See Levy and Markowitz (1979) for details.

If the individual is risk averse, $u''(\mu) < 0$, and expected utility decreases with variance. Thus, any expected utility function can be approximated as a function of mean and variance only. Using data on mutual fund returns, Levy and Markowitz (1979) show that these approximations are quite good for a broad class of utility functions. Regardless, a second-order approximation may be necessary since it is unlikely that resource managers will accurately perceive higher moments of the distribution of project returns or that data will accurately reflect their perceptions of these higher moments (Bockstael and Opaluch, 1983).

Thus, the manager chooses how to allocate the budget to maximize her expected utility:

$$\max_{w_1, w_2, \ldots, w_N} \int_{-\infty}^{\infty} U(\mu + \sigma z) f_Z(z) dz \tag{6}$$

subject to the restriction that the resulting policy portfolio must lie on the frontier: $\mu = Y(\sigma)$.[6] The welfare-maximizing policy portfolio $\boldsymbol{w}^* = [w_1^*, w_2^*, \ldots, w_N^*]$ is defined by the tangency between the policy frontier and the mean-variance indifference curve, shown in Fig. 1. At this point of tangency, society's marginal willingness to pay for reducing policy risk equals the price of reducing risk (i.e., society's marginal rate of substitution between risk and return equals the marginal rate of transformation between risk and return).

### 1.2. The welfare gains from diversified environmental policies

Fig. 1 illustrates how our mean-variance approach can be used to estimate the welfare gains from risk diversification. Because the indifference curve at $\boldsymbol{w}^*$ lies above the indifference curve with the entire budget allocated to the highest-return project (i.e., the portfolio at the far right on the frontier in Fig. 1), risk diversification generates welfare gains, measured as the difference in welfare implied by these two indifference curves. The welfare gains from risk diversification depend on the manager's willingness to pay for a risk reduction captured by the slope of the indifference curve. If the manager is not willing to pay for a risk reduction, the indifference curves in Fig. 1 would be flat, and the point of tangency would be on the far right of the frontier. In such a case there would be no welfare gains from risk diversification and the welfare-maximizing portfolio would only be composed of the high risk – high return project.

## 2. Estimating mean-variance indifference curves via choice experiments

As illustrated in the theoretical model in the previous section, the optimal policy portfolio and the welfare gains from risk diversification depend on the willingness to pay for a reduction in the variance associated with a portfolio of environmental projects. We propose using choice experiments to incorporate these risk preferences into portfolio theory. Choice experiments represent an attractive alternative to standard risk elicitation techniques such as Holt-Laury lottery menus which are typically used in the lab and incentivized to "control" for risk attitudes when analyzing the game of interest. If risk attitudes are context dependent (e.g., people are more willing to take on risk in a comfortable setting), the lab scenario may be a poor proxy for the habitat conservation decision of interest. While Holt-Laury lottery menus could be presented in a stated-preference survey to address concerns about context dependency, this raises two additional concerns. First, Holt and Laury (2002) compare risk elicitation in revealed and stated preference settings and find they diverge for large-stakes. Second, Holt-Laury experiments in a survey setting can be difficult to grasp since they force respondents to visualize the idea of making a series of choices but only having one of them 'count'. While randomization in a laboratory setting is straightforward, it is more challenging in a survey setting.

However, our theoretical model also reveals that standard empirical models employed in the choice experiment literature may undervalue risk diversification. Split sample studies have demonstrated that choice experiment respondents prefer certain outcomes to uncertain outcomes with the same expected value (e.g., Faccioli et al., 2019; Roberts et al., 2008; Rolfe and Windle, 2015). These studies typically present a respondent with an outcome, $x$, alongside $\pi$, the probability that $x$ will occur. Usually, $x$ and $\pi$ are separate attributes in the choice experiment, with their levels varying independently according to an experimental design. Modeling proceeds by specifying a functional form for utility based on assumptions about how respondents utilize and respond to the risk information (Glenk and Colombo, 2013; Wu et al., 2021). Most common is to assume respondents calculate the expected value of each alternative, which then enters the utility function linearly (i.e., the "mean only" model):

$$U_a = \beta_0 \mu_a + \varepsilon_a \tag{7}$$

where $U_a$ is the utility associated with alternative $a$, $\mu_a = \pi_a x_a$ is the expected value of attribute $x$ in alternative $a$, and $\beta_0$ is a parameter to be estimated representing the marginal utility of an increase in expected value. The common "mean only" empirical specification rules out risk aversion over attribute levels and implies flat rather than the upward sloping indifference curves in mean-variance space needed to generate welfare gains from risk diversification (see Fig. 1).

Another approach is to capture 'direct (dis)utility from risk' by estimating separate coefficients for $x_a$ and $\pi_a$. The direct disutility model often performs better in terms of model fit, but it lacks a theoretical foundation and thus has limited value for welfare analysis. A more limited number of studies have employed models allowing for nonlinear probability weighting by estimating the parameters of a weighting function (e.g., Roberts et al., 2008; Glenk and Colombo, 2013). However, it is not clear how to proceed with welfare analysis when using a probability weighting function (Holzer and McConnell, 2017). Regardless, neither of these empirical specifications can guarantee the upward sloping indifference curves in mean-variance space needed to generate welfare gains from risk diversification.

While variance in attribute levels will lower expected utility via the concavity of the utility function, linear utility specifications are inconsistent with risk aversion over income and commodity bundles (Hanoch, 1977; Holzer and McConnell, 2017). To illustrate, consider a choice experiment with standard expected utility framework and for which $x_1 > x_2$ represent the good and bad outcomes:

$$EU = \pi U(x_2) + (1-\pi) U(x_1) \tag{8}$$

$$Var[U] = \pi [U(x_2) - EU]^2 + (1-\pi)[U(x_1) - EU]^2 \tag{9}$$

While increases in the probability of the good state always increase expected utility, the relationship between probability and variance is nonmonotonic.[7] Thus, controlling for the expected value and probability in choice experiments may not capture individual risk preferences if

[6] Research on risk preferences in environmental contexts has shown that individuals often behave more in accordance with prospect theory than with expected utility theory (Riddel, 2012; Bartczak et al., 2015). In a financial setting, Levy and Levy (2004) find that one can employ mean-variance optimization to construct portfolios consistent with prospect theory. In an environmental application, Shah and Ando (2015) utilize downside risk measures to construct portfolios more consistent with behavior implied by portfolio theory.

[7] The relationship between the probability of the good state, 1-$\pi$, and the variance in utility is hump-shaped such that increases in the probability increase variance when that probability is low but decrease variance when that probability is high.

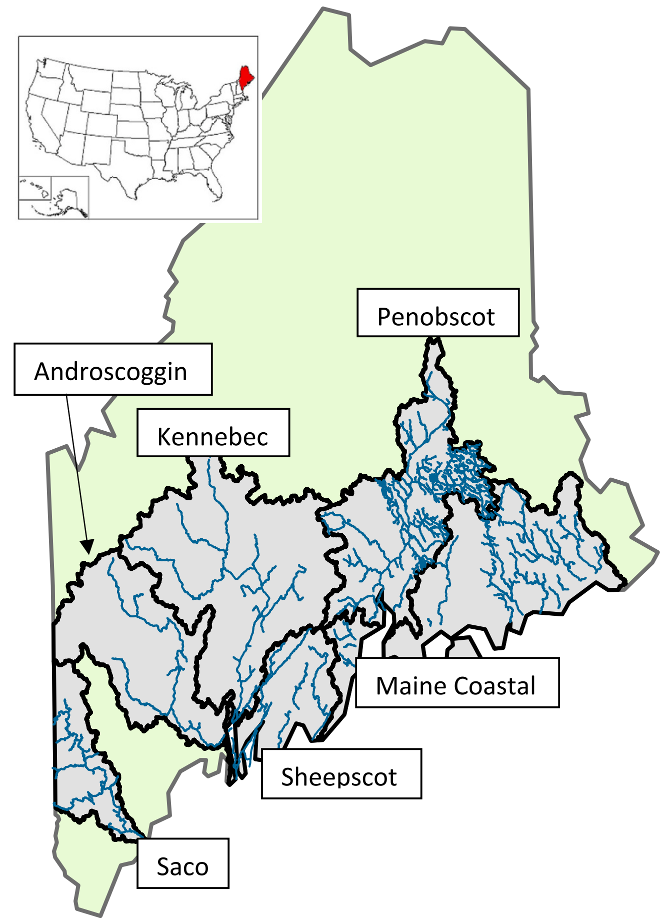


**Fig. 2.** Six watersheds in Maine are the potential projects (assets) in our restoration portfolio problem. Blue lines indicate the historic range of Atlantic salmon in each watershed. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

individuals are averse to variance in the policy outcomes.[8]

In contrast, a "mean-variance" empirical model would include variance as an additional explanatory variable. For example, consider the empirical model derived from the quadratic utility in Eq. (3):

$$U_a = \beta_1\mu_a + \beta_2\left(\sigma_a^2 + \mu_a^2\right) + \varepsilon_a \tag{10}$$

where $\sigma_a^2$ is the variance of $x$ in alternative $a$. Risk aversion implies that $\beta_2 < 0$. If $\beta_2 = 0$, implying risk neutrality, the mean-variance model collapses to the mean only model. If choice experiment respondents are risk averse, the mean only model in Eq. (7) is misspecified and resultant welfare estimates will be biased. To predict the direction of the bias, consider $\left(\sigma_a^2 + \mu_a^2\right)$ as an omitted variable from the mean only model in Eq. (7). We expect $\beta_2 < 0$ (i.e., risk aversion) which suggests the omitted variable will lead to an overestimate of welfare. However, positive correlation between $\mu_a$ and the omitted variable $\left(\sigma_a^2 + \mu_a^2\right)$ suggests that $\beta_0 < \beta_1$.[9]

The effect of omitted variable bias on welfare estimates is an empirical question, since it depends on the relative magnitudes of $\beta_0$ (from Eq. (7)), $\beta_1$ and $\beta_2$ (from Eq. (10)), and $\sigma_i^2$. Thus, the mean only model implies no welfare gains from risk diversification, encourages over investing in the most risky projects, and may overestimate the welfare from environmental projects if individuals are averse to the variance in project returns. To investigate the size of this bias and demonstrate the applicability of our mean-variance approach, we now turn to an application of our model to habitat restoration.

## 3. Application: habitat restoration in Maine

In December 2000, Gulf of Maine Atlantic salmon became officially listed as an endangered species. As a diadromous species, their lifecycle requires migration between fresh and saltwater habitat. The decline of Atlantic salmon is generally attributed to the extensive damming of Maine rivers, overfishing, climate change, and increased ocean mortality (Limburg and Waldman, 2009).

Recovery efforts focus on increasing abundance of salmon through a system of conservation hatcheries and enhancing survivability by improving river habitat connectivity and quality via dam removal (NOAA, 2019). Due to the multiple pressures facing Atlantic salmon, and the poor understanding of ocean mortality drivers, recovery efforts have proven challenging and results cannot be guaranteed (Mills et al., 2013). At the same time, recovery is expensive. Recovery costs are estimated at $24 million annually, of which just $8.6 million is currently funded (NOAA, 2019). The combination of recovery expense, budget shortfall, and outcome uncertainty provide an ideal context for our analysis of preferences for environmental risk and return.

### 3.1. Ecological data and the portfolio frontier

The first step in deriving the portfolio frontier is to define the restoration projects and quantify the return measure. We define projects as watersheds in Maine that are currently or were historically utilized by Atlantic salmon for spawning habitat. Investing in conservation and/or restoration of each watershed will have a return (expected future salmon population) and risk (variance of the expected outcome).[10] To quantify risk and return, we begin with annual counts of salmon in each watershed, then scale them based on historic habitat extent to generate a plausible time series of salmon in each watershed under full restoration. From the time series, we obtain the three necessary inputs to derive a portfolio frontier - mean and variance for each watershed, as well as the covariance between watersheds.

We obtained annual fish counts from the U.S. Atlantic Salmon Assessment Committee Annual Report (NOAA, 2020). The counts originate from trapping operations run by state agencies and hydropower companies throughout New England rivers. Atlantic salmon counts are available for 2010–2019, and the complete time series exists for ten rivers. We aggregate counts by watershed (HUC8 level), yielding six projects (Fig. 2). The counts provide a reasonable proxy for the effect of multiple underlying risks on variation in the population, as well as the relationships between the impact of those risks in different watersheds (i.e., covariance between projects). However, they do not reflect the

[8] A risk premium (expected value minus the certainty equivalent) is consistent with variance as the appropriate measure of risk. However, a risk premium captures the WTP for eliminating risk and is therefore of little use when investigating changes in risk.

[9] Although $\mu$ and $\sigma^2$ are both determined by $x$ and $\pi$, the non-monotonic relationship between $\mu$ and $\sigma^2$ is such that they can be considered independent here. The positive correlation between $\mu$ and the omitted variable $(\sigma^2 + \mu^2)$ arises solely from the mathematical relationship between $\mu$ and $\mu^2$.

[10] By defining return as the expected future salmon population, we implicitly assume that the cost of salmon restoration in each watershed is the same. This assumption may be accurate for restoration activities like hatchery development and stocking efforts which are likely similar across watersheds. However, this assumption may not hold for dam removal and modification which will depend on the unique dam configuration in each watershed. As more dam removals take place, estimates of the cost of dam removal in each watershed will improve allowing our biological return measure to be converted to a return on investment (ROI) that measures the expected future salmon per dollar invested in dam removal.

restoration potential of each watershed.

To reflect restoration potential more closely, we scaled the annual counts based on the extent of historic Atlantic salmon habitat in each watershed. We utilized salmon habitat data obtained from an extensive effort to map the historic range of diadromous fish in Maine (Houston et al., 2007). We began the scaling with an estimate for the maximum number of Atlantic salmon that could be supported by full restoration of historic habitat in the Penobscot watershed (8000) from Blachly et al. (2023). We used the mean level for 'likelihood that salmon will utilize newly restored habitat' from the choice experiment (63 %) to estimate an expected salmon count for the Penobscot under complete restoration ($8,000 \times 0.63 = 5,040$). Next, we divided by the length of historic habitat in the Penobscot to obtain an estimate for 'expected salmon per unit of habitat'. Then, we multiplied by the amount of habitat in each watershed to obtain watershed-level expected salmon counts. This approach implies the expected salmon per unit habitat in each watershed is based on the Penobscot watershed; the location of the most well-documented river restoration project in Maine.[11] Scale factors for each watershed were derived by dividing the expected salmon counts by the mean values observed in the annual counts. The scale factors were then applied to each observation in the original time series.

Expected counts and standard deviations for each watershed in the first two columns of Table 2 provide a reasonable estimate of the combined effect of multiple sources of uncertainty in salmon populations under full restoration. For example, populations post restoration may be uncertain due to uncertainty in salmon fishing effort, natural variability in spawning habitat or ocean conditions, or the inability to predict whether salmon will utilize restored habitat. The Penobscot is a high risk, high return watershed while the Sheepscot is a relatively low risk watershed.

We used the Portfolio function in MATLAB to perform the optimization in Eq. (1). The resulting frontier identifies allocations (weights) of the habitat restoration budget across watersheds that maximize the expected number of salmon for each level of risk (variance of the expected outcome).

The frontier is presented in Fig. 3 (top). At one extreme, the high risk/return portfolio (point (a)) is achieved by allocating 100 % of the restoration budget to the single highest risk/return asset (the Penobscot watershed). It is possible to achieve lower variance (i.e., move to the left along the frontier) by shifting the restoration budget away from the Penobscot. The slope of the frontier indicates the amount of expected return that must be foregone to achieve a marginal reduction in variance. In this way, one can think of the frontier as quantifying the price of reducing risk, which is dictated by covariance relationships within the ecological system. The concave shape of the frontier indicates that the price of a risk reduction increases as variance decreases until both risk and return are minimized (point (b)) by spreading the restoration budget across all six watersheds. If society is extremely risk averse (i.e., willingness to pay for reducing risk is very high) welfare will be maximized at point (b). On the other hand, if society is risk neutral (i.e., willingness to pay for the risk reduction is 0) welfare will be maximized at point (a). Between those extremes, some other budget allocation is optimal.

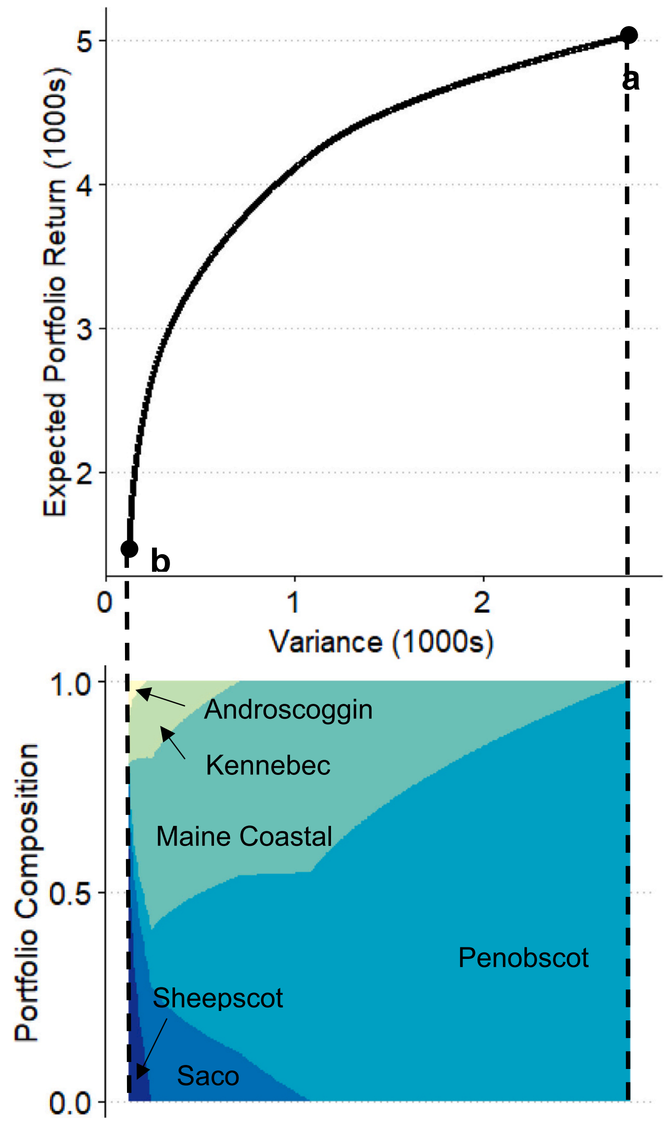


**Fig. 3.** Portfolio frontier (top panel) and the composition of each policy portfolio along the frontier (bottom panel). The portfolio that maximizes return (point a) represents 100 % allocation to the Penobscot. The portfolio that minimizes risk (point b) splits the budget across all six watersheds.

### 3.2. Preference data and choice modeling

To identify the welfare maximizing point on the frontier in Fig. 3, we utilized data from a choice experiment designed to elicit public preferences for ecosystem service tradeoffs provided by dam removals, upgrades, and modifications in Maine. Over the course of three focus groups we determined the four most relevant ecosystem service endpoints: Atlantic salmon habitat, river herring habitat, hydropower production, and recreation. We utilized fish habitat and hydrologic models to determine a feasible range for each attribute, chose discrete levels within each range, and used those levels to create a D-efficient design (Rose and Bliemer, 2009). Full design and implementation details can be found in the appendix and Blachly et al. (2023). In this paper we focus entirely on the risky attribute—Atlantic salmon.

Since Atlantic salmon recovery operations are inherently risky, presenting outcomes as certain would likely produce biased welfare estimates (Johnston et al., 2017). In addition to an attribute describing the number of salmon that can be supported by restored habitat, therefore, we introduced an additional attribute "likelihood that salmon will utilize newly restored habitat" (see Fig. 4). This is in line with the most common approach to outcome uncertainty in the literature, where a single outcome and its probability vary exogenously as attributes.[12] To utilize the mean-variance modeling approach, we calculated the

[11] Completed in 2016, the Penobscot River Restoration Project involved decommissioning and removing two lower mainstem dams in exchange for increasing hydro capacity at an upstream dam.

[12] However, Makriyannis et al. (2024) find that percentage probabilities may not provide useful information to stated preference respondents beyond that conveyed by past event frequencies.

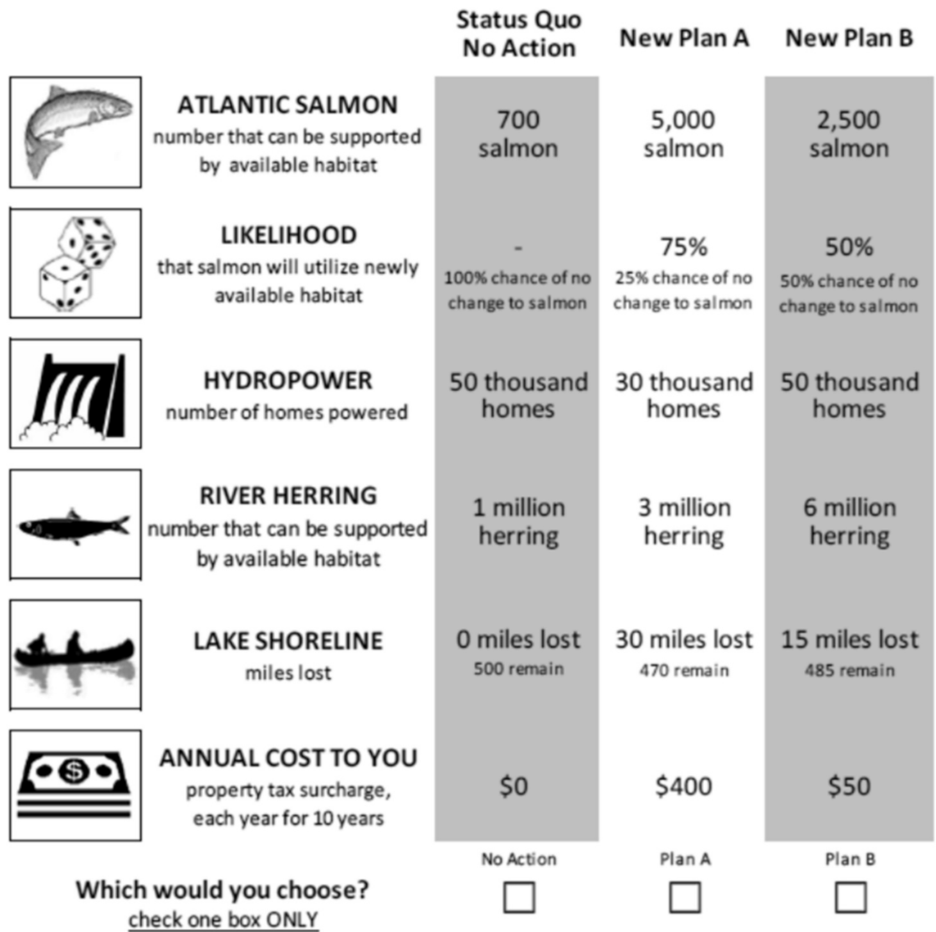


**Fig. 4.** A sample choice set. The salmon and likelihood attributes combine to represent discrete distributions that can be described by a mean $\mu_{salm}$ and variance $\sigma^2_{salm}$.

expected value and variance of the salmon outcome implied by each choice alternative.

We analyzed the choice experiment data within the standard random utility framework. We assume that respondents maximize their utility, so the probability that a particular respondent chooses alternative $a$ is: $P_a = Prob[U_a > U_b \forall b \neq a]$. Further, we assume that utility contains observable and random components that are additively separable ($U_a = v_a + \varepsilon_a$), and that the random components are independent and identically distributed following an Extreme Value Type I distribution. This yields the familiar multinomial logit choice model (McFadden, 1973). We utilize two specifications for $v_a$ ("mean only" and "mean-variance"), representing different assumptions about the risk preferences of respondents. Our baseline ("mean only") model replicates the most common approach to modeling uncertain choice experiment outcomes:

$$v_a = \beta_{SQ} + \beta_{salm}\pi_a x_{salm\,a} + \beta_{hyd} x_{hyd\,a} + \beta_{her} x_{her\,a} + \beta_{lake} x_{lake\,a} + \beta_{cost} x_{cost\,a} \tag{11}$$

It assumes that respondents use the provided information about the carrying capacity of newly restored salmon habitat, $x_{salm}$, and the likelihood that salmon will utilize the new habitat, $\pi$, to form an expectation of the outcome described by alternative $a$. To simplify comparisons with the mean-variance specification, we drop the $a$ subscripts, write the expectation as $\mu_{salm}$ ($= \pi \cdot x_{salm}$), and collapse the remaining $k$ attributes into a summation, so that Eq. (11) becomes: $U = \beta_{salm}\mu_{salm} + \sum_k \beta_k x_k + \varepsilon$.

Our mean-variance model requires an assumption about the specification of the underlying utility function. The constant absolute risk aversion utility function $u(x_{salm}) = 1 - \exp(-\alpha x_{salm})$ requires that $x_{salm}$ be normally distributed. To avoid imposing assumptions about the distribution of $x_{salm}$, we follow the quadratic utility specification in Eq. (3)[13]:

$$v = \beta_{salm}\mu_{salm} + \beta_{\sigma^2}\left(\sigma^2_{salm} + \mu^2_{salm}\right) + \sum_k \beta_k x_k \tag{12}$$

where $\mu_{salm} = \pi \cdot x_{salm}$ is the expected outcome of the choice alternative and $\sigma^2_{salm} = \mu_{salm} \cdot (1-\pi)$ is the associated variance of the choice alternative. We estimate Eqs. (11) and (12) using mixed (random parameter) logit to account for the panel structure of the data and allow for individual taste heterogeneity. All parameters are assumed to follow normal distributions, except for $\beta_{cost}$, which is held fixed. We perform likelihood ratio tests to determine if the models are significantly different in terms of fit.

Random parameter logit results for the two utility specifications are reported in Table 1. All coefficients are statistically significant and have the expected sign. Respondents derive positive utility from increasing habitat for Atlantic salmon (row 1) and river herring (row 4), while maintaining hydropower capacity (row 3) and avoiding reductions to lake shoreline (row 5). Generally, respondents prefer some new program to maintaining the status quo (row 6). The mean-variance specification

[13] Quadratic utility $u(x_{salm}) = x_{salm} - b{x_{salm}}^2$ implies utility declines when $x_{salm} > 1/(2b)$. As we show in Fig. 6, this is not a major concern for our application since the vast majority of the probability density is in the region $x < 1/(2b)$.

**Table 1**
Mixed logit results.

| | Mean only | Mean-variance |
|---|---|---|
| | Means of random parameters | |
| $\mu_{\text{salmon}}$ | 0.149*** | 1.285*** |
| | (0.421) | (0.239) |
| $\mu^2, \sigma^2_{\text{salmon}}$ | | −0.129*** |
| | | (0.026) |
| Hydropower | 0.013*** | 0.013*** |
| | (0.003) | (0.003) |
| Herring | 0.068*** | 0.085*** |
| | (0.021) | (0.023) |
| Lakes | −2.042*** | −3.240*** |
| | (0.472) | (0.565) |
| Status Quo | −1.772*** | −1.258*** |
| | (0.421) | (0.473) |
| | Standard deviations | |
| $\mu_{\text{salmon}}$ | 0.143** | 0.061 |
| | (0.071) | (0.081) |
| $\mu^2, \sigma^2_{\text{salmon}}$ | | 0.001 |
| | | (0.012) |
| Hydropower | 0.013*** | 0.012** |
| | (0.004) | (0.005) |
| Herring | 0.013 | 0.008 |
| | (0.030) | (0.032) |
| Lakes | 2.400*** | 2.587*** |
| | (0.541) | (0.516) |
| Status Quo | 4.903*** | 4.924*** |
| | (0.591) | (0.600) |
| | Fixed parameters | |
| Cost | −0.006*** | −0.006*** |
| | (0.001) | (0.001) |
| LL | −767.3 | −755.2 |
| AIC | 1556.6 | 1536.4 |
| Pseudo-$R^2$ | 0.3916 | 0.4012 |

Notes: $N = 1148$ observations across 197 individuals. Random parameters follow normal distributions.

(column 2) indicates diminishing marginal utility for salmon along with distaste for variance (row 2). Statistically significant standard deviation estimates for the random parameters indicate heterogeneous preferences for hydropower, lake shoreline, and maintaining the status quo. On the other hand, insignificant standard deviation coefficients suggest preferences for river herring are more consistent across the sample. The preference heterogeneity for salmon present in column 1 disappears when we allow for risk aversion in column 2. One possible explanation is omitted variable bias in the mean-only model.

Results suggest that implied outcome variance is a significant driver of preferences. A likelihood ratio test confirms that the mean-variance specification (column 2) fits the data significantly better than the more common mean-only model in column 1 ($p$-value = 0.00). Several studies have measured the disutility associated with outcome variance (uncertainty) compared to zero variance (certainty) (e.g., Faccioli et al., 2019), but to our knowledge model 2 is the first to treat variance as a continuous attribute, allowing an estimate of marginal disutility. The model implies marginal willingness to pay of \$21.50 to decrease the variance of restoration outcomes by 1000 salmon, holding the mean level of salmon constant.

It is also informative to analyze preferences for the tradeoff between risk and expected return (i.e., risk preferences) as indifference curves in risk-return space. Fig. 5 shows the indifference curves implied from the responses to our choice experiment. The mean-only model in column 1 is linear in the risky attribute (salmon), implying risk neutrality and yielding flat indifference curves (Fig. 5 – left pane). The mean-only model assumes respondents hold no preferences over variance, so utility is strictly increasing in expected value, irrespective of variance. In contrast, the mean-variance model in column 2 produces upward sloping, convex indifference curves (Fig. 5 – right pane), implying risk aversion.[14] Respondents will accept higher variance in the outcome of the restoration policy but only if they are compensated in the form of higher expected salmon populations. We turn next to the welfare implications of analyzing mean and variance jointly.

### 3.3. Welfare implications

To generate welfare estimates for uncertain salmon restoration scenarios in Fig. 4 we follow Hanemann (1984):

$$CS = -\frac{1}{\beta_{cost}}\left[\ln\sum_n \exp v_n^1 - \ln\sum_n \exp v_n^0\right] \tag{13}$$

where CS is the compensating surplus welfare measure and $v_n^0$ and $v_n^1$ are the $n$th respondent's indirect utility before and after restoration, respectively. For these calculations, we hold levels of the other attributes constant between $v^0$ and $v^1$ to focus on the impact of salmon risk and return. Results suggest that the mean-only model underestimates the average respondent's welfare associated with uncertain salmon outcomes compared to the mean-variance specification. Columns 3 and 5 in Table 2 show that the underestimation of welfare by the mean-only

[14] The mean-variance indifference curve for the absolute risk aversion utility function $u(x) = k - \exp(-\alpha x)$ is $\mu = \frac{\alpha}{2}\sigma^2 - \frac{\ln(1-k)}{\alpha}$. Thus, the upward sloping indifference curves produced from the quadratic utility function are consistent with the absolute risk aversion utility function with a positive risk aversion coefficient $\alpha > 0$.

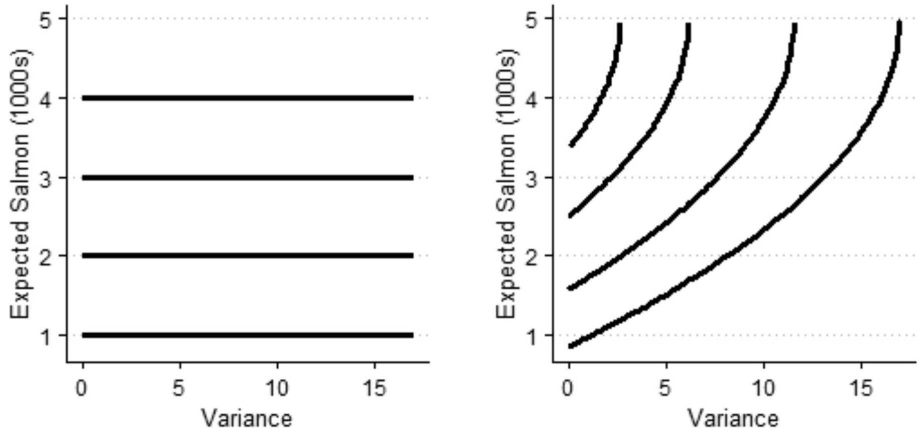


**Fig. 5.** Mean-variance indifference curves derived from responses to our choice experiment using the mean-only model (left) and the mean-variance model (right).

**Table 2**
Expected return, risk (SD), and welfare estimates for individual watersheds and the optimal policy portfolio.

| | | | Mean only model | | Mean-variance model | | |
|---|---|---|---|---|---|---|---|
| HUC 8 Watershed | E [Salmon] | SD | Average welfare ($/person) | Optimal restoration budget allocation | Average welfare ($/person) | Optimal restoration budget allocation | Simple diversification budget allocation |
| Androscoggin | 1039 | 1335 | 25.80 | 0 | 160.99 | 0 | 1/6 |
| Kennebec | 1881 | 1563 | 46.71 | 0 | 274.25 | 0 | 1/6 |
| Maine Coastal | 3231 | 1187 | 80.24 | 0 | 437.23 | 0.39 | 1/6 |
| Penobscot | 5040 | 1679 | 125.16 | 1 | 472.66 | 0.61 | 1/6 |
| Saco | 1177 | 1486 | 29.23 | 0 | 174.81 | 0 | 1/6 |
| Sheepscot | 1243 | 541 | 30.87 | 0 | 226.70 | 0 | 1/6 |
| Portfolio E[Salmon] | | | | 5040 | | 4326 | 2268 |
| Portfolio SD | | | | 1679 | | 1109 | 526 |
| Portfolio Welfare | | | | 125.16 | | 497.70 | 363.88 |

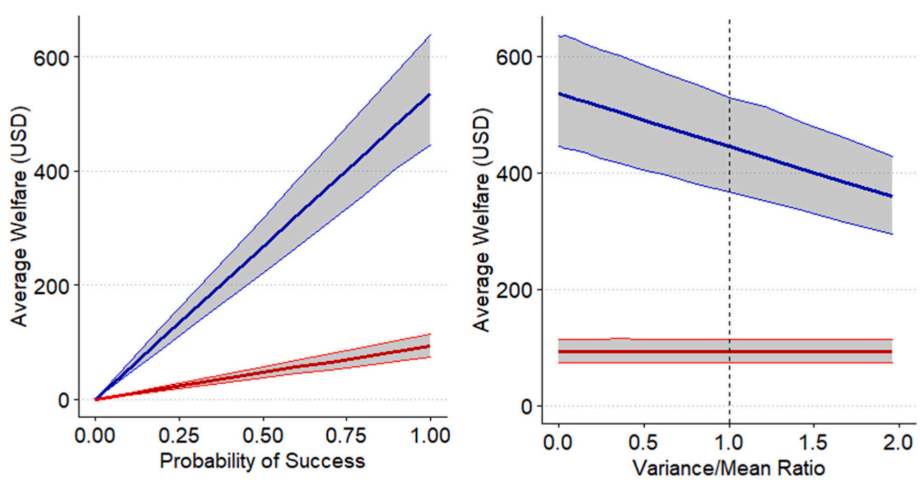


**Fig. 6.** Average welfare of (left) restoring habitat to support 4000 salmon with increasing probability of success; and (right) restoring habitat expected to support 4000 salmon holding probability of success at 0.50 while increasing variance via a mean-preserving spread under the mean-only (red lines) and mean-variance (blue lines) utility models. Gray shaded areas indicate 95 % confidence intervals. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

model occurs in each of the six Maine watersheds and varies from \$160.99–\$25.80 = \$135.19 (Androscoggin) to \$437.23–\$80.24 = \$356.99 (Maine Coastal).

Consider a hypothetical habitat restoration project to support 4000 additional salmon with an unknown probability of success (Fig. 6 – left pane). As probability of success increases from 0 % to 50 %, expected value and variance increase. From 50 % to 100 %, however, expected value continues to increase while variance decreases. The mean-only model (red lines) underestimates welfare compared to the mean-variance model (blue lines) across the full range of probabilities, especially at the high end where expected value is high but variance is low. As illustrated in our theoretical model, the mean-only model suffers from omitted variable bias. Failing to control for changes in salmon variance depresses the marginal effect of salmon means resulting in a downwardly biased estimate for the coefficient on $\mu_{salm}$. This effect is larger than the effect of the negative coefficient on $\mu_{salm}^2 + \sigma_{salm}^2$. When respondents are averse to variance, a mean-only model that fails to control for that variance will underestimate welfare.

The mean-variance model also implies that it is theoretically possible for a scenario with positive expected value to produce negative welfare if variance is sufficiently large. Thus, the observed underestimation of welfare produced by the mean-only model may dissipate when outcome variance is high. Consider another situation where habitat restoration is expected to support 4000 additional salmon with increasing variance (i.e., a mean-preserving spread) (Fig. 6 – right pane). The mean-only model produces a constant welfare estimate regardless of variance, while estimates from the mean-variance specification decrease with variance due to risk aversion. Results show, however, that the mean-only model produces significantly lower welfare estimates even when salmon returns are over-dispersed (variance/mean ratio $> 1$). We conclude that the mean-only model will underestimate welfare for all reasonably dispersed salmon outcomes.

### 3.4. Policy portfolio choice and welfare gains from diversification

We utilize our empirical measure of risk preferences (i.e., mean-variance indifference curves derived from the choice experiment) to identify the welfare-maximizing policy portfolio (Table 2 Columns 4 and 6). Results indicate that society is willing to sacrifice some expected return to achieve lower risk. In particular, welfare is maximized by splitting the restoration budget between the Penobscot (61 %) and Maine Coastal (39 %) watersheds. The allocation results in expected annual salmon of 4326 with a standard deviation of 1109. At this point, marginal willingness to pay for reducing risk equals the price of the risk reduction dictated by the ecological system (i.e., the tangency condition is met) resulting in average welfare per person of $497.70.

In contrast, the flat indifference curves implied by the mean-only model suggest zero willingness to pay to reduce risk, which results in 100 % of the restoration budget allocated to the Penobscot. Thus, omitting variance when estimating welfare for uncertain restoration policies has two effects: 1) it underestimates welfare associated with restoration in each watershed and 2) it redirects the restoration budget to the riskiest watershed. Comparing the optimal portfolio based on the mean-variance welfare estimates ($497.70) and the optimal portfolio for the mean-only welfare estimates ($125.16) provides an estimate of the overall bias that results from both of these effects.

The portfolio choice results indicate that social welfare is improved by shifting restoration away from the watershed with the highest expected return (Penobscot) to the watershed with lower risk (Maine Coastal). Because salmon returns to the two watersheds are not perfectly correlated, splitting the restoration budget results in lower overall risk than could be achieved from investing in either individually. The welfare gains from a diversified salmon restoration policy are found by comparing the mean-variance welfare estimates for the optimal portfolio ($497.70) and the Penobscot-only portfolio ($472.66). Specifically, about 5 % of the welfare from salmon restoration ($25.04 = $497.70–$472.66) can be attributed to the value society holds for reducing the riskiness of salmon restoration.

The low value of risk diversification may be due to our application to salmon restoration. Because salmon spend most of their lives in the ocean, environmental variation that affects populations in one watershed will likely affect them all. The result is a high covariance in salmon populations across watersheds. To investigate the implications of this aspect of salmon lifecycles, we lower each element in the covariance matrix by 50 % and calculate new optimal portfolio weights. The uniform reduction in covariances would mimic a case where more environmental shocks take place in the rivers and less in the ocean. The results, provided in the appendix, suggest less but more valuable diversification. The lower covariance reduces the standard deviation associated with the portfolio weights in Table 2. Given the risk preferences implied from the indifference curve, the resulting optimal portfolio places slightly more of the budget in the Penobscot and increases the portfolio welfare from $497.70 to $512.36. The welfare gains from a diversified salmon restoration policy increase to $39.70 (nearly 8 % of the welfare from salmon restoration). Because of the shared risks of anadromous species, we expect the value of diversification to be higher in other applications to habitat restoration.

Our welfare maximizing portfolio can also be used to estimate the welfare losses from employing a common diversification rule of thumb. Simple diversification evenly divides the budget across all assets to avoid the risk associated with investing in a single asset. A simple diversification strategy may also be adopted to promote equity if taxpayers oppose having federal conservation dollars focused on a single area. This simple diversification point rarely lies on the efficient frontier, resulting in efficiency losses often measured as the decrease in expected benefits for a given level of risk (Ando and Mallory, 2012). Table 2 shows the welfare losses when employing this simple diversification strategy. When the budget is evenly divided across all six watersheds (portfolio weights equal 1/6), the risk is much lower than the optimal portfolio. However, this lower risk comes at the expense of lower expected return resulting in a lower average welfare per person of $363.88. Welfare is 37 % lower when using this simple diversification strategy. In our salmon restoration example, no diversification (i.e., investing only in the Penobscot) results in higher welfare than simple diversification cautioning against conservation strategies that overly diversify risk.

## 4. Conclusions

Environmental policy with limited budgets often requires choosing between many environmental projects. We present a framework for determining the welfare-maximizing funding allocation when the outcomes of these environmental projects are uncertain. We use the framework to show how welfare gains from a diversified environmental policy depend on society's willingness to pay for reducing risk and the price of risk reduction, determined by the covariance between projects. To estimate the willingness to pay, we propose a mean-variance utility approach to analyze stated preference data that accounts for aversion to outcome variability and leverages a popular choice experiment format.

In an application to salmon restoration in Maine, we show that optimal diversification increases welfare by 5 % relative to no diversification. However, a commonly used diversification rule-of-thumb results in excessive diversification which lowers welfare by 37 %. While we present an application of our model to habitat restoration, our general approach can be used to identify welfare-maximizing budget allocations in other settings such as uncertain pollution control across multiple point sources and catch limits for multiple fisheries.

This study demonstrates the usefulness of a mean-variance model for choice experiment applications to risky environmental projects. Our mean-variance approach, which simplifies risky decisions to functions of the mean and variance only, may be especially useful for public land managers and non-profit managers who do not know the utility functions of their constituents. In these settings, mean-variance approximations can provide investment strategies that almost maximize expected utility (Levy and Markowitz, 1979). However, our results are subject to several caveats which provide a roadmap for future work. First, the data we utilized did not come from a choice experiment designed for mean-variance analysis, necessitating strong assumptions about how respondents inferred variance from the simplified risk information provided. Future work should study the credibility of this assumption beyond our test of model fit. This could be done by building variance into the experimental design, and split-samples could be used to test the relative effectiveness of different methods for conveying variance as both an implied and explicit attribute. Possible approaches for building variance explicitly into the experimental design have been explored in transportation (Li et al., 2010), flood risk (Makriyannis et al., 2018), and species conservation (Gerling et al., 2024) contexts.

In addition, we note issues that arose from applying portfolio theory in a restoration setting. Unlike traditional finance applications where the returns to individual assets are exogenous, restoration implies endogeneity (Loreau et al., 2021). That is, allocation of restoration activity (dam removal) to a particular asset (watershed) changes the expected return (number of salmon) to that asset. In contrast, existing applications to large-scale conservation (e.g., Ando and Mallory, 2012; Vinent et al., 2019) are more akin to the finance case, where the portfolio return from conservation equals the current return weighted by the portfolio allocation. We chose to utilize the fully restored watersheds as our return measure, implying that any allocation of restoration would yield the fully restored expected outcome. More work is needed to incorporate ecological production functions for each asset into the risk-return optimization, reflecting the fact that individual asset returns are a function of portfolio weights. Finally, the welfare gains from diversification warrant further research, particularly to develop theory and generate evidence in a variety of settings.

## Acknowledgements and disclaimers

We are grateful to Christian Vossler, Amy Ando, and seminar participants at the University of Wyoming, The Arctic University of Norway, and Oregon State University for helpful comments and suggestions on an earlier version of the paper. The findings and conclusions in this report are those of the authors and should not be construed to represent any official USDA or U.S. Government determination or policy. This research was supported [in part] by the U.S. Department of Agriculture, Forest Service. Any use of trade, firm, or product names is for descriptive purposes only and does not imply endorsement by the U.S. government. Sims gratefully acknowledges grant support from the USDA National Institute of Food and Agriculture, Grant/Award Numbers RI0018-W4133 and 11121636.

## CRediT authorship contribution statement

**Ben Blachly:** Writing – review & editing, Writing – original draft, Visualization, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Charles Sims:** Writing – review & editing, Writing – original draft, Supervision, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization. **Travis Warziniack:** Writing – review & editing, Investigation, Funding acquisition, Conceptualization.

## Declaration of competing interest

The author declares that (s)he has no relevant or material financial interests that relate to the research described in this paper.

## Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.ecolecon.2025.108750.

## Data availability

Data will be made available on request.

## References

Ando, A.W., Mallory, M.L., 2012. Optimal portfolio design to reduce climate-related conservation uncertainty in the prairie pothole region. Proc. Natl. Acad. Sci. 109, 6484–6489. https://doi.org/10.1073/pnas.1114653109.

Ando, A., Fraterrigo, J., Guntenspergen, G., Howlader, A., Mallory, M., Olker, J.H., Stickley, S., 2018. When portfolio theory can help environmental investment planning to reduce climate risk to future environmental outcomes—and when it cannot. Conserv. Lett. 11, e12596. https://doi.org/10.1111/conl.12596.

Bartczak, A., Chilton, S., Meyerhoff, J., 2015. Wildfires in Poland: the impact of risk preferences and loss aversion on environmental choices. Ecol. Econ. 116, 300–309. https://doi.org/10.1016/j.ecolecon.2015.05.006.

Blachly, B., Uchida, E., Roy, S.G., 2023. Integrating public preferences with biophysical production possibilities: an application to ecosystem services from dam removal. Ecol. Soc. 28, 1. https://doi.org/10.5751/ES-13739-280151.

Bockstael, N.E., Opaluch, J.J., 1983. Discrete modelling of supply response under uncertainty: the case of the fishery. J. Environ. Econ. Manag. 10 (2), 125–137.

Crowe, K.A., Parker, W.H., 2008. Using portfolio theory to guide reforestation and restoration under climate change scenarios. Clim. Chang. 89, 355–370. https://doi.org/10.1007/s10584-007-9373-x.

de Palma, A., et al., 2008. Risk, uncertainty, and discrete choice models. Mark. Lett. 19, 269–285.

Faccioli, M., Kuhfuss, L., Czajkowski, M., 2019. Stated preferences for conservation policies under uncertainty: insights on the effect of individuals' risk attitudes in the environmental domain. Environ. Resour. Econ. 73, 627–659. https://doi.org/10.1007/s10640-018-0276-2.

Gerling, C., Bartczak, A., Hanley, N., 2024. Accounting for Risk in Species Conservation Programmes. Available at SSRN:. https://doi.org/10.2139/ssrn.4760451.

Glenk, K., Colombo, S., 2013. Modelling outcome-related risk in choice experiments. Aust. J. Agric. Resour. Econ. 57, 559–578. https://doi.org/10.1111/1467-8489.12012.

Halpern, B.S., White, C., Lester, S.E., Costello, C., Gaines, S.D., 2011. Using portfolio theory to assess tradeoffs between return from natural capital and social equity across space. Biol. Conserv. 144 (5), 1499–1507. https://doi.org/10.1016/j.biocon.2011.01.019.

Hanemann, W.M., 1984. Welfare evaluations in contingent valuation experiments with discrete responses. Am. J. Agric. Econ. 66, 332–341. https://doi.org/10.2307/1240800.

Hanoch, G., 1977. Risk aversion and consumer preferences. Econometrica 45 (2), 413–426.

Holt, C.A., Laury, S.K., 2002. Risk aversion and incentive effects. Am. Econ. Rev. 92 (5), 1644–1655.

Holzer, J., McConnell, K., 2017. Risk preferences and compliance in recreational fisheries. J. Assoc. Environ. Resour. Econ. 4 (S1), S1–S35.

Houston, R., Chadbourne, K., Lary, S., Charry, B., 2007. Geographic Distribution of Diadromous Fish in Maine. U.S. Fish and Wildlife Service, Gulf of Maine Coastal Program, Falmouth, Maine.

Johnston, R.J., Boyle, K.J., Adamowicz, W., Bennett, J., Brouwer, R., Cameron, T.A., Hanemann, W.M., Hanley, N., Ryan, M., Scarpa, R., Tourangeau, R., Vossler, C.A., 2017. Contemporary guidance for stated preference studies. J. Assoc. Environ. Resour. Econ. 4, 319–405. https://doi.org/10.1086/691697.

Leroux, A.D., Martin, V., 2016. Hedging supply risks: An optimal urban water portfolio. Am. J. Agric. Econ. 98 (1), 276–296.

Leroux, A.D., Martin, V.L., Zheng, H., 2018. Addressing water shortages by force of habit. Resour. Energy Econ. 53, 42–61.

Leroux, A.D., Martin, V.L., St. John, K. A., 2022. Modeling time varying risk of natural resource assets: implications of climate change. Quant. Econ. 13 (1), 225–257.

Levy, H., Levy, M., 2004. Prospect theory and mean-variance analysis. Rev. Financ. Stud. 17 (4), 1015–1041.

Levy, H., Markowitz, H.M., 1979. Approximating expected utility by a function of mean and variance. Am. Econ. Rev. 69 (3), 308–317.

Li, Z., Hensher, D.A., Rose, J.M., 2010. Willingness to pay for travel time reliability in passenger transport: A review and some new empirical evidence. Transp. Res. Part E 46 (3), 384–403.

Limburg, K.E., Waldman, J.R., 2009. Dramatic declines in North Atlantic diadromous fishes. BioScience 59, 955–965. https://doi.org/10.1525/bio.2009.59.11.7.

Loreau, M., Barbier, M., Filotas, E., Gravel, D., Isbell, F., Miller, S.J., Montoya, J.M., Wang, S., Aussenac, R., Germain, R., Thompson, P.L., Gonzalez, A., Dee, L.E., 2021. Biodiversity as insurance: from concept to measurement and application. Biol. Rev. 96, 2333–2354. https://doi.org/10.1111/brv.12756.

Makriyannis, C., Johnston, R.J., Whelchel, A.W., 2018. Are choice experiment treatments of outcome uncertainty sufficient? An application to climate risk reductions. J. Agric. Resour. Econ. 1–33. https://doi.org/10.1017/age.2017.27.

Markowitz, H., 1952. Portfolio selection. J. Financ. 7.

Matthies, B.D., Jacobsen, J.B., Knoke, T., Paul, C., Valsta, L., 2019. Utilising portfolio theory in environmental research – new perspectives and considerations. J. Environ. Manag. 231, 926–939. https://doi.org/10.1016/j.jenvman.2018.10.049.

McFadden, D., 1973. Conditional Logit Analysis of Qualitative Choice Behavior. Academic Press, New York.

Meles, T.H., Lokina, R., Mtenga, E.L., Tibanywana, J.J., 2023. Stated preferences with survey consequentiality and outcome uncertainty: A Split sample discrete choice experiment. Environ. Resour. Econ. 86 (4), 1–38.

Merton, R.C., 1969. Lifetime portfolio selection under uncertainty. Rev. Econ. Stat. 51 (3), 247–257.

Merton, R.C., 1971. Optimum consumption and portfolio rules in a continuous-time model. J. Econ. Theory 3 (4), 373–413.

Mills, K.E., Pershing, A.J., Sheehan, T.F., Mountain, D., 2013. Climate and ecosystem linkages explain widespread declines in north American Atlantic salmon populations. Glob. Chang. Biol. 19, 3046–3061. https://doi.org/10.1111/gcb.12298.

NOAA, 2019. Recovery Plan for the Gulf of Maine Distinct Population Segment of Atlantic Salmon. National Marine Fisheries Service; US Fish and Wildlife Service.

NOAA, 2020. Annual Report of the U.S. Atlantic Salmon Assessment Committee (No. 32). Portland, Maine.

Pindyck, R.S., 2007. Uncertainty in environmental economics. Rev. Environ. Econ. Policy 1, 45–65. https://doi.org/10.1093/reep/rem002.

Pindyck, R.S., 2014. Risk and return in the Design of Environmental Policy. J. Assoc. Environ. Resour. Econ. 1, 395–418. https://doi.org/10.1086/677949.

Riddel, M., 2012. Comparing risk preferences over financial and environmental lotteries. J. Risk Uncertain. 45, 135–157. https://doi.org/10.1007/s11166-012-9149-1. New York.

Roberts, D.C., Boyer, T.A., Lusk, J.L., 2008. Preferences for environmental quality under uncertainty. Ecol. Econ. 66, 584–593. https://doi.org/10.1016/j.ecolecon.2008.05.010.

Rolfe, J., Windle, J., 2015. Do respondents adjust their expected utility in the presence of an outcome certainty attribute in a choice experiment? Environ. Resour. Econ. 60, 125–142. https://doi.org/10.1007/s10640-014-9760-5.

Rose, J.M., Bliemer, M.C.J., 2009. Constructing efficient stated choice experimental designs. Transp. Rev. 29, 587–617. https://doi.org/10.1080/01441640902827623.

Shah, P., Ando, A.W., 2015. Downside versus symmetric measures of uncertainty in natural resource portfolio design to manage climate change uncertainty. Land Econ. 91 (4), 664–687.

Sharpe, W.F., 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. *j*. Financ. 19 (3), 425–442.

Tobin, J., 1958. Liquidity preference as behavior towards risk. Rev. Econ. Stud. 25, 65–86.

Vinent, O.D., Johnston, R.J., Kirwan, M.L., Leroux, A.D., Martin, V.L., 2019. Coastal dynamics and adaptation to uncertain sea level rise: optimal portfolios for salt marsh migration. J. Environ. Econ. Manag. 98, 102262. https://doi.org/10.1016/j.jeem.2019.102262.

von Eiff, D., Pommeret, A., 2021. Determining an optimal water mix for Hong Kong. Environ. Model. Assess. 26 (2), 221–236.

von Eiff, D., Yeo, J., An, A.K., Chopra, S.S., 2023. Comparative economic and life cycle analysis of future water supply mix scenarios for Hong Kong–A water Scarce City. J. Environ. Manag. 325, 116370.

Wu, H., Mentzakis, E., Schaafsma, M., 2021. Exploring different assumptions about outcome-related risk perceptions in discrete choice experiments. Environ. Resour. Econ. https://doi.org/10.1007/s10640-021-00638-x.