

QUANTITATIVE FINANCE
RESEARCH CENTRE



UNIVERSITY OF
TECHNOLOGY SYDNEY



QUANTITATIVE FINANCE RESEARCH CENTRE

Research Paper 393

September 2018

Model Risk Measurement under Wasserstein Distance

Yu Feng and Erik Schlögl

ISSN 1441-8010

www.qfrc.uts.edu.au

Model Risk Measurement under Wasserstein Distance

Yu Feng¹ and Erik Schlögl¹

¹University of Technology Sydney, Quantitative Finance Research Centre

September 11, 2018

Abstract

The paper proposes a new approach to model risk measurement based on the Wasserstein distance between two probability measures. It formulates the theoretical motivation resulting from the interpretation of fictitious adversary of robust risk management. The proposed approach accounts for all alternative models and incorporates the economic reality of the fictitious adversary. It provides practically feasible results that overcome the restriction and the integrability issue imposed by the nominal model. The Wasserstein approach suits for all types of model risk problems, ranging from the single-asset hedging risk problem to the multi-asset allocation problem. The robust capital allocation line, accounting for the correlation risk, is not achievable with other non-parametric approaches.

1 Introduction

The current work on robust risk management either focuses on parameter uncertainty or relies on comparison between models. To go beyond that, Glasserman and Xu recently proposed a non-parametric approach. Under this framework, a worst-case model is found among all alternative models in a neighborhood of the nominal model. Glasserman and Xu adopted the Kullback-Leibler divergence (i.e. relative entropy) to measure the distance between an alternative model and the nominal model. They also proposed the use of the α -divergence to avoid heavy tails that causes integrability issues under the Kullback-Leibler divergence.

Both the Kullback-Leibler divergence and the α -divergence are special examples of the f -divergence [1, 2, 3]. A big problem of using f -divergence is that it is well-defined only when the alternative measure is absolutely continuous with respect to the nominal measure. This limits the range of the alternative models under consideration. In some cases, we may want to search over all possible probability measures, whether they are absolutely continuous or not. This is especially true when we apply this approach to volatility, which corresponds to the quadratic variation of a process. If the process is driven by a Brownian motion, then searching over absolutely continuous measures rules out any volatility risk. In Fig. 1(a), the distribution of the volatility is a Dirac- δ function under the nominal model. The worst-case scenario that accounts for the volatility risk has a widely spread distribution of the volatility. However, f -divergence is not well-defined in this case, and therefore the worst-case scenario simply gets ignored.




Fig 1: (a) Dirac measure has a support of a single point. The alternative model with a widespread distribution cannot be achieved using f -divergence. (b) State transition in a metric space. f -divergence does not involves the metric, so the transition from State 1 to 2 takes the same amount of cost as the transition from 1 to 3. Remember to redraw the diagram to lift the base line from 0 to 0.2

Furthermore, the state space considered by financial practitioners is usually equipped by a natural metric. For instance, the price of a security takes value from the set of positive real numbers, and thus naturally inherits the Euclidean metric. Assuming a diffusion process, the price of the security moves along a continuous path. This means that a large price change is less probable than a small price change, implying a stronger deviation from the nominal model. However, the distance of the move, measured by the natural metric, is not explicitly taken into account when using f -divergence. Fig. 1(b) shows three models corresponding to three distributions of the security price. Assuming the Model 1 is adopted as the nominal model, then Model 2 as an alternative model is apparently more probable than Model 3. However, one cannot tell the difference using any type of f -divergence.

As an attempt to solve these issues, we suggest to adopt the Wasserstein metric to measure the distance between probability measures. Relying on the metric equipped in the state space, the Wasserstein metric works for any two measures, even if their supports are mutually exclusive. As a result, the proposed Wasserstein approach accounts for all alternative measures instead of merely the absolutely continuous ones. These features allow us to resolve the two issues of the f -divergence as mentioned above. For financial practitioners, the proposed approach is especially useful when dealing with nominal measures with a subspace support (such as a Dirac measure). This paper is organized in the following manner. Sec. 2.1 offers a conceptual introduction including the intuitive motivation and the basics about the Wasserstein metric and its associated transportation theory. Sec. 3.1 is the theoretical part that provides the problem formulation and main results. It also includes practical considerations and comparison between different approaches. Sec. 4 gives a few interesting applications in mathematical finance, ranging from the volatility risk in option pricing and hedging to robust portfolio optimisation.

2 Basic Concepts

2.1 Motivation and Adversary Interpretation

To illustrate the idea of model risk in an intuitive way, we start from a simple discrete-state space. An example is the credit rating which is ordinal, e.g. A+, A, A-, BBB+,

etc. Assuming we have a nominal model that states that in a month the credit rating of an institution could be A+, A- or BBB+. The nominal model assigns probabilities of 25%, 50% and 25% to the three states. Since we do not possess complete information, model risk exists either because the actual probabilities of the three states are different or because other ratings are still possible. Glasserman and Xu proposed the so-called “adversary” interpretation which suggests a fictitious adversary that perturbs the probabilities against us [4]. By perturbing the probabilities essentially the adversary adds new information, limited by its information entropy budget. For example, if the adversary would like to move 5% chance from A+ to BBB+, its consumption of relative entropy is

$$0.2 \ln \left(\frac{0.2}{0.25} \right) + 0.3 \ln \left(\frac{0.3}{0.25} \right) = 0.01 \quad (1)$$

Now suppose the adversary would like to move the 5% chance to BBB, which is a state of 0 probability under the nominal measure. The consumption of relative entropy

$$0.2 \ln \left(\frac{0.2}{0.25} \right) + 0.05 \ln \left(\frac{0.05}{0} \right) \quad (2)$$

becomes infinite. This simply means that such perturbation is impossible no matter how much control the adversary has. In the language of probability theory, relative entropy is well-defined only when the new measure is absolutely continuous with respect to the nominal one.

To allow for a more generic quantification of model risk, we may re-define the requested cost of perturbation. Instead of using the relative entropy, we consider about the cost of a state transition (termed as the transportation cost). This transportation cost is usually given by some metric equipped by the state space (to ensure its symmetry). For simplicity we assume that the distance between two credit rating is given by the number of ratings in between, e.g. $d(A+, A-)=1$ and $d(A+, BBB+)=2$. We calculate the weighted average transportation costs for the two types of perturbations discussed in the last paragraph:

1. shift 5% chance from A+ to BBB+: transportation cost= $5\% \times 2=0.1$
2. shift 5% chance from A+ to BBB: transportation cost= $5\% \times 3=0.15$

The second-type perturbation only involves a cost slighter larger than the first type, instead of being infinite.

Using the transportation cost described above, one is capable of measuring the adversary’s cost for all alternative measures rather than merely the absolutely continuous ones. However, it may provide state transitions that are over-concentrated, thus not reflecting the actual competitive market structure. To illustrate this point, think about the transition from state A+. The fictitious adversary, assumed to be one agent, would push the rating only in one direction. This implies that the transportation performed by the fictitious agent can be represented by a (deterministic) map on the state space $T : \Omega \rightarrow \Omega$. T is referred as the **transportation map** [5]. In fact, suppose it is optimal to the fictitious agent (thus the worst case scenario) to transit the state A+ to a state, say BBB+. There is no motivation for the agent to transport any density on A+ to other states. This results from the linearity of the transportation cost, and will be illustrated further in the next section.

Glasserman and Xu’s interpretation of model risk involves an fictitious adversary but without explicit consideration of its economic nature. They assume that the adversary perform uniformly aiming to maximise our expected loss. In reality, such an adversary can only be achieved by a single agent or institution. The actual market structure, however, is usually more competitive. In economic terms, the fictitious adversary may consist of

heterogeneous agents who act independently. This asks for approaches that quantify the model risk based on the actual market structure.

Now get back to the credit rating example. In reality there might be multiple agents that are capable of impacting the rating, among which some prefer to upgrade the rating while others prefer to downgrade the rating. This asks for a different formulation of state transitions, for the final state transitioned from a given initial state becomes a random variable. All we know is a probability measure conditional to the given initial state (or a transition density). Overall, the transportation is described by a joint probability density $\gamma : \Omega \times \Omega \rightarrow \mathbb{R}^+$ instead of a deterministic map. The joint density (or the corresponding measure on $\Omega \times \Omega$) is referred as the **transportation plan**[6]. This allows us to formulate the optimisation problem w.r.t the transportation plan instead of the transportation map. Such formulation leads to more general results capable of accounting for different types of market structure.

From a practical perspective, the main advantage of using the Wasserstein metric is to deal with nominal measures supported by strict subspaces. Still in the example of credit rating, the nominal measure is supported by $\{A+, A-, BBB+\}$, which is a strict subspace of the entire state space (of rating). Approaches based on f -divergence are only capable of incorporating alternative measures with the same support. Using the Wasserstein approach, on the other hand, does allow us to alter the support. In particular, if we formulate the problem using a transportation map T , then the new support is $\{T(A+), T(A-), T(BBB+)\}$, still a strict subspace. Therefore, although different transportation maps provide us with different supports, none of them is capable of spreading to the entire state space. On the other hand, by formulating the problem with a transportation plan, we indeed account for alternative measures that are supported by the entire space. Now regarding the fictitious adversary as a class of heterogeneous agents, it is reasonable to believe that the distribution is widely spread under the perturbation of the adversary.

In conclusion, we are interested in an approach to model risk measurement that formulates the transportation cost based on a transportation plan. We will see that this approach is capable of account for actual market structure by parametrising an entropy constraint (Sec. 3.2). In the remaining part of this section, we will review the Wasserstein metric and its associated transportation theory.

2.2 Transportation Theory and Wasserstein Metric

Starting from this point, we will always assume a continuous-state space unless otherwise stated. The approach for discrete-state spaces follows the same routine and therefore is omitted. Now let the state space (Ω, d) be a Polish metric space, we may define the transportation cost $c : \Omega \times \Omega \rightarrow \mathbb{R}_+$ by the n -th power of the metric, i.e. $c(x, y) = d(x, y)^n$, where $n \in [1, \infty)$. Given two probability measures P and Q on (Ω, d) , we may formulate the optimal transportation problem using either a transportation map or a transportation plan. For the former approach, we aim to find the transportation map $T : \Omega \rightarrow \Omega$ that realizes the infimum

$$\begin{aligned} \inf_T \int_{\Omega} p(x) c(x, T(x)) dx & \quad (3) \\ \text{s.t. } |J_T(x)| q(T(x)) = p(x), \forall x \in \Omega & \end{aligned}$$

where $p(x)$ and $q(x)$ are the probability density functions of the two measures P and Q , respectively. J_T is the Jacobian of the map T . It is part of the constraint that enforces the map T to be measure-preserving. Eq. 3 is referred as the **Monge's formulation** of the optimal transportation problem.

The problem of Monge's formulation is that the existence of a measure-preserving map T is not guaranteed. Examples in the last section provide a discrete-state illustration of this issue: $\text{supp}(Q) = \{T(A+), T(A-), T(BBB+)\}$ has at most three elements. As a result, there is no measure-preserving map if $|\text{supp}(Q)| > |\text{supp}(P)|$. In a continuous-state space, a measure-preserving map sends a Dirac measure to another Dirac measure. Therefore, measure-preserving map does not exist if P is a Dirac measure while Q is not. The ill-posed Monge's formulation can be improved by adopting a transportation plan $\gamma : \Omega \times \Omega \rightarrow \mathbb{R}_+$:

$$\begin{aligned} & \inf_{\gamma} \int_{\Omega \times \Omega} \gamma(x, y) c(x, y) dx dy & (4) \\ & \text{s.t.} \int_{\Omega} \gamma(x, y) dy = p(x) \\ & \int_{\Omega} \gamma(x, y) dx = q(y) \end{aligned}$$

Eq. 4 is referred as the **Kantorovich's formulation** of the optimal transportation problem. It is clear that every transportation map T can be given by a transportation plan

$$\gamma(x, y) = |J_T(x)| q(y) \delta(y - T(x)) \quad (5)$$

where $\delta(\cdot)$ is the Dirac- δ function. In addition, the existence of a transportation plan is guaranteed as $\gamma(x, y) = p(x)q(y)$ always satisfies the constraints in Eq. 4. According to these observations, the Kantorovich's formulation is preferred over the Monge's formulation. Remember that the transportation cost $c(x, y)$ is the n -th power of the metric $c(x, y)$. The n -th **Wasserstein matrix**, denoted by W_n , is defined as the infimum in Eq. 4, raised to the power of $1/n$. In the next section, the theoretical formulation and the main results of this paper will be presented with the help of the Kantorovich's formulation. The transportation cost function $c(x, y)$ will be regarded as a generic non-negative function, without reference to its specific form or the power n .

3 Theory

3.1 Wasserstein Formulation of the Model Risk Problem

The core part of model risk measurement is to determine the alternative model under the worst-case scenario. In the language of probability theory, we need to determine the alternative probability measure that maximizes our expected loss. We may formulate the problem in the following way. Given a nominal probability measure P on the state space Ω , we would like to find a worst-case measure Q^* that realizes the following superimum:

$$\begin{aligned} & \sup_Q \mathbb{E}^Q[V(X)] & (6) \\ & \text{s.t.} D(P||Q) \leq \eta \end{aligned}$$

The expectation is taken under the alternative measure Q , on a loss function $V : \Omega \rightarrow \mathbb{R}$. Only alternative measures that are close enough to the nominal measure are deemed as legitimate. This restriction is formulated by constraining the statistical distance $D(P||Q)$ to be equal to or less than a constant η .

Glasserman and Xu suggest using the relative entropy (or Kullback-Leibler divergence) for $D(P, Q)$. Like any f -divergence, relative entropy has limited feasibility as only equiv-

alent measures are legitimate. Based on the discussion in the last section, we suggest to apply the Wasserstein metric instead. The actual formulation of the model risk problem, on the other hand, has a slightly different form than Eq. 6. Specifically, instead of optimizing the expectation w.r.t the alternative measure Q (or its density function $q : \Omega \rightarrow \mathbb{R}^+$), we optimize the expectation w.r.t the transportation plan $\gamma : \Omega \times \Omega \rightarrow \mathbb{R}^+$ directly. The single constraint on q is replaced by two constraints applied to γ , including the marginalisation condition given in Eq. 4. This formulation roots from the idea of state transition and is illustrated as below.

Based on the discussion in the last section, for any pair of states $x, y \in \Omega$ all we need to find is the transition density from x to y , $p_{Y|X}(y|x)$. Given a function of transportation cost from x to y , $c(x, y)$, the expected transportation cost conditional to an initial state x is

$$W(x) = \int_{\Omega} p_{Y|X}(y|x)c(x, y)dy \quad (7)$$

The initial state x follows a distribution $p_X(x)$ given by the nominal model. Take expectation under the nominal measure, we get the unconditional transportation cost

$$W = \int_{\Omega} p_X(x)W(x)dx = \int_{\Omega \times \Omega} p_{X,Y}(x, y)c(x, y)dxdy \quad (8)$$

where the joint distribution $p_{X,Y}(x, y) = p_X(x)p_{Y|X}(y|x)$. To be consistent with the notation used previously, we denote the marginal distributions p_X, p_Y by p, q , and the joint distribution $p_{X,Y}$ by the transportation plan γ . It is noted that the transition converts the initial distribution $p(x)$ to a final distribution $q(y)$, inducing a change of measure on the state space Ω .

One of the key tasks of the model risk measurement is to solve for the worst-case model under certain constraints. These constraints set the criteria for legitimate alternative models. Now denote the loss function by $V(x)$ ($x \in \Omega$), the probability density function of the nominal model by $p(x)$, and the probability density function of an alternative model by $q(x)$. We formulate the problem by the superimum of the expected loss over all legitimate models:

$$\sup_{q(y)} \int_{\Omega} q(y)V(y)dy \quad (9)$$

According to the discussion in the last section, we regard the change of measure as probabilistic state transitions. The probability density function $q(y)$ of the alternative model is merely the marginalisation of a joint density (or transportation plan) $\gamma(x, y)$, i.e. $q(y) = \int_{\Omega} \gamma(x, y)dx$. This allows us to take superimum over $\gamma(x, y)$ instead of $q(y)$:

$$\sup_{\gamma(x,y)} \int_{\Omega \times \Omega} \gamma(x, y)V(y)dxdy \quad (10)$$

The first constraint of the superimum problem comes from the marginalisation of the joint density w.r.t x , as it is given by the nominal model:

$$\int_{\Omega} \gamma(x, y)dy = p(x) \quad (11)$$

In a similar way to Glasserman and Xu's work, we restrict all alternative measures by their distances from the nominal model. The distance is now measured by the average transportation cost given in Eq. 8. It reflects the expected cost paid by a fictitious adversary

who attempts to transit a state x to an alternative state y according to the transportation plan $\gamma(x, y)$. This results in the following constraint which defines the set of legitimate measures:

$$\int_{\Omega \times \Omega} \gamma(x, y) c(x, y) dx dy \leq \eta \quad (12)$$

The constant η in Eq. 12 is termed as the Wasserstein distance budget, just as the relative entropy budget in Glasserman and Xu's approach [4]. In order to account for a specific density function $q^*(y)$ in the constrained supermum problem given by Eq. 9-12, the Wasserstein distance, defined in Eq. 4, between $q^*(y)$ and the nominal density $p(x)$ cannot exceed η . In fact, if $q^*(y)$ can be obtained by marginalizing a transportation cost $\gamma^*(x, y)$ that satisfies Eq. 11-12, then according to Eq. 4 its Wasserstein distance with the nominal density function $p(x)$ is

$$\begin{aligned} W(p, q^*) &= \inf_{\gamma} \int_{\Omega \times \Omega} \gamma(x, y) c(x, y) dx dy \\ &\leq \int_{\Omega \times \Omega} \gamma^*(x, y) c(x, y) dx dy \leq \eta \end{aligned} \quad (13)$$

On the other hand, if $W(p, q^*) < \eta$, then the density function $q^*(y)$ can always be expressed by the marginalisation of a transportation plan $\gamma^*(x, y)$ that satisfies Eq. 11-12. Otherwise, in the definition of the Wasserstein distance, Eq. 4, η sets a lower bound for the term

$$\int_{\Omega \times \Omega} \gamma(x, y) c(x, y) dx dy \quad (14)$$

Therefore the Wasserstein distance, as the infimum of the term above, is equal to or larger than η . This immediately violates the assumption $W(p, q^*) < \eta$. In summary, η sets the maximum level (budget) of Wasserstein distance for an alternative measure to be legitimate.

Remarkably, even though the problem (Eq. 10-12) is formulated using the transportation plan (Kantorovich's formulation), its solution can be expressed by a transportation map $T^* : \Omega \rightarrow \Omega$,

$$T^*(x) = \arg \max_{y \in \Omega} \left[V(y) - \frac{c(x, y)}{\beta} \right] \quad (15)$$

where $\beta \in \mathbb{R}_{++}$ is a constant. The underlying reason is the linearity of Eq. 14 w.r.t the transportation plan γ . Suppose the worst case scenario is to transit a state x to another state $T^*(x)$. Then there is no motivation for the fictitious adversary to transit x to states other than $T^*(x)$, say $T'(x)$, for the adversary could continue improving the target by increasing $\gamma(x, T^*(x))$ while reducing $\gamma(x, T'(x))$ (by the same amount). See Appendix A for a sketch of the derivation of Eq. 15.

3.2 Entropy Constraint on Transportation Plan

Eq. 15 provides the worst-case transportation map for the problem formulated in Eq. 10-12. This formulation in fact assumes a zero-sum game between two parties, in which our counterparty attempts to shift a state x to $y^*(x)$ (deterministically) so that its profit (thus our loss) can be maximized. In Sec. 2.1, we mentioned that the actual market structure may be more competitive consisting of heterogeneous agents that act more or less independently.

This calls for a widespread transition density $p_{Y|X}(y|x)$ (instead of being a δ -function).

In practice, it is also advantageous of having a widely distributed transition density. For the purpose of risk management, we need to consider a wide range of alternative measures due to model ambiguity. As a result a widespread distribution is usually more representative than a narrow distribution. From the information-theoretic point of view, a widespread distribution contains less information (more entropy) thus more appropriately representing the model ambiguity. Now have a think about the practical situations where the approaches based on f -divergence are not applicable. They usually have nominal measures that are too restrictive in the sense that they are supported by merely subspaces (of the state space). To correctly quantify the model risk one should consider widespread distributions supported by the entire state space. However, these distributions do not have well-defined f -divergence w.r.t the nominal measure, providing an inherent issue of these approaches.

One of the primary purposes of using Wasserstein metric instead of f -divergence is to tackle this issue. Specifically, we would like to include all measures regardless of their support. This purpose is achieved by using the Kantorovich's formulation as illustrated in Sec. 2.2. However, without further constraint the worst-case model can still be achieved with a transportation map, as illustrated by Eq. 15. This causes the worst-case measure to be restrictive if the nominal measure is supported by merely a subspace. To achieve a widespread worst-case distribution, one may need to impose further constraints to Eq. 10-12.

A Dirac nominal measure, denoted by P , provides a special example where Eq. 15 is not suitable of characterizing the worst-case scenario. Applying the transportation map T^* results in the worst-case measure supported by $\{T(x)\}$ where x is the sole element in $supp(P)$. The worst-case measure is Dirac as well. In most cases, this worst-case measure inappropriately accounts for model ambiguity. To resolve this issue, we may further impose an entropy constraint that guarantees the worst-case measure to be supported by the entire state space:

$$-\int_{\Omega \times \Omega} \gamma(x, y) \ln \gamma(x, y) dx dy \geq \mu \quad (16)$$

The LHS is the (differential) entropy [7] of the joint distribution (transportation plan) $\gamma(x, y)$, and the RHS is a constant $\mu \in \mathbb{R}$ (or a positive constant $\mu \in \mathbb{R}_{++}$ for discrete-state space). This constraint excludes every transportation plan that is equivalent to a transportation map. In fact, every transportation map T gives a transportation plan with a δ -function transition density (see Eq. 5). For such transportation plan, the δ -function makes the LHS of Eq. 16 approaching negative infinity (or zero for discrete-state space), and is therefore excluded.

Alternatively, Eq. 16 can be interpreted with respect to the transition density function $p_{Y|X}(y|x)$. We may rewrite Eq. 16 by

$$-\int_{\Omega \times \Omega} \gamma(x, y) \ln p_{Y|X}(y|x) dx dy \geq \mu + \int_{\Omega} p(x) \ln p(x) dx \quad (17)$$

Eq. 17 imposes restriction to the transition density function. A tighter restriction (with a larger μ) implies a wider transition density, reflecting a market structure that is more competitive. On the other hand, if we relax the constraint completely by shifting μ towards negative infinity (or zero for discrete-state space), then we permit transition densities to take the form of δ -functions, corresponding to the single-agent adversary.

We may further introduce terms from information theory, and rewrite Eq. 17 by

$$\int_{\Omega} p(x)H(Y|X = x)dx \geq \mu - H(X) \quad (18)$$

where $H(X)$ denotes the entropy of the random variable X [7]. Since its distribution $p(x)$ is given by the nominal model, $H(X)$ is deemed as a constant. $H(Y|X = x)$, on the other hand, is the information entropy w.r.t the transition density $p_{Y|X}(y|x)$. It is interpreted as the entropy of the random variable Y , conditional to X taking a given value x . $H(Y|X = x)$ quantifies the uncertainty of the transportation from a given state x . Generally a more competitive market that involves more independent decision-makers leads to a more uncertain state transition, thus a larger $H(Y|X = x)$. As a result, Eq. 18 allows us to incorporate the actual market structure by parametrising μ . It is noted that in information theory, the LHS of Eq. 19 is termed as the conditional (differential) entropy and is denoted by $H(Y|X)$ [7]. This leads to an equivalent information-theoretic version of the constraint Eq. 16:

$$H(Y|X) \geq \mu - H(X) \quad (19)$$

3.3 Main Result and Discussion

The supremum problem Eq. 10 subject to the three constraints Eq. 11, 12 and 16 formulates the complete version of the Wasserstein approach to model risk measurement. Now suppose there exists a joint distribution $\gamma^*(x, y)$ that solves the problem. Then the worst-case model is characterised by a probability density function

$$q^*(y) = \int_{x \in \Omega} \gamma^*(x, y)dx, \quad \forall y \in \Omega \quad (20)$$

To solve the constrained supremum problem, we introduce two multipliers $\alpha \in \mathbb{R}_+$ and $\beta \in \mathbb{R}_+$, and transform the original problem to a dual problem. Solving the inner part of the dual problem leads to our main result (see Appendix B for derivation):

$$q^*(y) = \int_{\Omega} dx \frac{p(x) \exp\left(\frac{V(y)}{\alpha} - \frac{c(x,y)}{\alpha\beta}\right)}{\int_{\Omega} \exp\left(\frac{V(y)}{\alpha} - \frac{c(x,y)}{\alpha\beta}\right) dy} \quad (21)$$

It is noted that the multipliers α and β are in fact controlling variables that determines the levels of restriction, of the entropy constraint Eq. 16 and the transportation constraint Eq. 12, respectively.

The limit when α approaches zero corresponds to complete relaxation of the entropy constraint Eq. 16. In this limit Eq. 20 degenerates to the probability density function induced by the transportation map given by Eq. 15. On the other side of the spectrum, Eq. 20 gives a uniform distribution when α approaches infinity, as a result of the tight entropy constraint.

In the extreme case of $\beta = 0$, Eq. 20 leads to a simple result $q^*(x) = p(x)$. This is because the transportation constraint Eq. 12 reaches its tightest limit ($\eta = 0$). No state transition is allowed thus preserving the nominal model. On the other hand, when β approaches infinity, the worst-case distribution $q^*(y) \sim \exp(V(y)/\alpha)$ is exponentially distributed. In this case, the transportation cost is essentially zero. As a result, the worst-case measure is the one that maximises the expected value of $V(Y)$ with a reasonably large entropy (the maximum expected value is given by a Dirac measure at $\arg \max_y V(y)$ but this results in

a very low entropy). Special cases of Eq. 21 are tabulated in Tab. 1 for different values of α and β .

Table 1: Worst-case probability density function at different (α, β) combinations. p is the nominal distribution and u is the uniform distribution. δ denotes the Dirac δ -function and T^* is the transportation map given by Eq. 15.

	$\alpha = 0$	α	$\alpha \rightarrow \infty$
$\beta = 0$	$p(x)$		
β	$p(T^{*-1}(x))/ J_T $	given by Eq. 20	$u(x)$
$\beta \rightarrow \infty$	$\delta(x - \arg \max V(x))$	$\propto e^{V(x)/\alpha}$	

3.4 Practical Consideration

According to Table. 1, the worst-case measure follows a uniform distribution when α approaches infinity (i.e. under the most restrictive entropy constraint). In practice, we may want the worst-case distribution converges to a given density function q_0 instead of being uniform. This requires modification on the formulation of the problem, by generalising the entropy constraint Eq. 19 into

$$-D_{KL}(P(Y|X)||Q_0(Y)) \geq \mu - H(X) - H(Y) \quad (22)$$

$D_{KL}(P(Y|X)||Q_0(Y))$ denotes the conditional relative entropy, given by the expected value of the KL divergence, $D_{KL}(P(Y|X = x)||Q_0(Y))$, of the two probability density functions w.r.t y , $p_{Y|X}(\cdot|x)$ and $q_0(\cdot)$. Written explicitly, the conditional relative entropy takes the form of

$$\begin{aligned} D_{KL}(P(Y|X)||Q_0(Y)) &= \int_{\Omega} p(x) \left(\int_{\Omega} p_{Y|X}(y|x) \ln \left(\frac{p_{Y|X}(y|x)}{q_0(y)} \right) dy \right) dx \\ &= \int_{\Omega \times \Omega} \gamma(x, y) \ln \frac{\gamma(x, y)}{q_0(y)} dx dy - \int_{\Omega} p(x) \ln p(x) dx \end{aligned} \quad (23)$$

Substituting Eq. 23 into Eq. 22 allows us to obtain the explicit version of the constraint:

$$- \int_{\Omega \times \Omega} \gamma(x, y) \ln \frac{\gamma(x, y)}{q_0(y)} dx dy - \int_{\Omega} q_0(y) \ln q_0(y) dy \geq \mu \quad (24)$$

It is clear that the previous entropy constraint Eq. 16 is merely a special case of Eq. 24 in which q_0 is a uniform distribution. Under this formulation, the problem that we need to solve consists of Eq. 10, 11, 12 and 24. The result differs from Eq. 21 by a weighting function q_0 (see Appendix B for derivation):

$$q^*(y) = \int_{\Omega} dx \frac{p(x)q_0(y) \exp \left(\frac{V(y)}{\alpha} - \frac{c(x,y)}{\alpha\beta} \right)}{\int_{\Omega} q_0(y) \exp \left(\frac{V(y)}{\alpha} - \frac{c(x,y)}{\alpha\beta} \right) dy} \quad (25)$$

It is noted that Eq. 25 takes a similar form to the Bayes' theorem and q_0 serves as the

prior distribution. In fact, if the conditional distribution takes the following form:

$$p_{X|Y}^*(x|y) \propto \exp\left(\frac{V(y)}{\alpha} - \frac{c(x,y)}{\alpha\beta}\right) \quad (26)$$

Then the Bayes' theorem states that

$$\begin{aligned} p_{Y|X}^*(y|x) &= \frac{p_{X|Y}^*(x|y)q_0(y)}{\mathbf{E}_Y\left(p_{X|Y}^*(x|\cdot)q_0(\cdot)\right)} \\ &= \frac{q_0(y) \exp\left(\frac{V(y)}{\alpha} - \frac{c(x,y)}{\alpha\beta}\right)}{\int_{\Omega} q_0(y) \exp\left(\frac{V(y)}{\alpha} - \frac{c(x,y)}{\alpha\beta}\right) dy} \end{aligned} \quad (27)$$

which is the posterior distribution of Y given the observation $X = x$. Now if we observe a distribution $p(x)$ over X , then we may infer the distribution of Y to be

$$\begin{aligned} q^*(y) &= \int_{\Omega} p(x)p_{Y|X}^*(y|x)dx \\ &= \int_{\Omega} dx \frac{p(x)q_0(y) \exp\left(\frac{V(y)}{\alpha} - \frac{c(x,y)}{\alpha\beta}\right)}{\int_{\Omega} q_0(y) \exp\left(\frac{V(y)}{\alpha} - \frac{c(x,y)}{\alpha\beta}\right) dy} \end{aligned} \quad (28)$$

which is exactly the worst-case distribution given in Eq. 25.

The connection between the Bayes' theorem and Eq. 25 is not just a coincident. In fact, the worst-case distribution of Y , given in Eq. 28, can be regarded as the posterior distribution of a latent variable. On the other hand, the nominal model of X , given by $p(x)$, is considered as the distribution that is actually observed. Assuming no nominal model exists (i.e. no observation on X has been made), then our best guess on the latent variable Y is given solely by its prior distribution $q_0(y)$. Now if the observable variable X does take a particular value x , then we need to update our estimation according to the Bayes' theorem (Eq. 27). The conditional probability density $p_{X|Y}^*(x|y)$ takes the form of Eq. 26, reflecting the fact that the observable variable X and the latent variable Y are not far apart. Imaging that we generate a sampling set $\{x_i\}$ following the nominal distribution $p(x)$, then for each x_i we get a posterior distribution $p_{Y|X}^*(y|x_i)$ from Eq. 27. Overall, the best estimation of the distribution over the latent variable Y results from the aggregation of these posterior distributions. This is achieved by averaging them weighted by their probabilities $p(x_i)$, as given in Eq. 28. This leads to the Bayesian interpretation of the model risk measurement, which concludes that by "observing" the nominal model $p(x)$ over the observable variable X , the worst-case model is given by updating the distribution of the latent variable Y , from the prior distribution $q_0(y)$ to the posterior distribution $q^*(y)$.

If we know nothing about the nominal model, setting the prior q_0 to a uniform distribution seems to make the most sense (because a uniform distribution maximizes the entropy thus containing least information). This leads to the main result given by Eq. 21. However, it is sometimes much more convenient to choose a prior other than the uniform distribution. A particular interesting case is to set q_0 the same as the nominal distribution p . In this case, the limit of $\beta \rightarrow \infty$ (complete relaxation of the transportation constraint) is given by

$$q^*(x) = \frac{p(x)e^{\theta V(x)}}{\int_{\Omega} p(x)e^{\theta V(x)} dx} \quad (29)$$

Table 2: Worst-case density function with prior q_0 at different (α, β) combinations. p is the nominal distribution. δ denotes the Dirac δ -function and T^* is the transportation map given by Eq. 15.

	$\alpha = 0$	α	$\alpha \rightarrow \infty$
$\beta = 0$	$p(x)$		
β	$p(T^{*-1}(x))/ J_T $	given by Eq. 25	$q_0(x)$
$\beta \rightarrow \infty$	$\delta(x - \arg \max V(x))$	$\propto q_0(x)e^{V(x)/\alpha}$	

where we replace the parameter α^{-1} by θ . This limit is exactly the worst-case distribution given by the relative entropy approach [4]. Despite of the simplicity of Eq. 29, it is not recommended to set $q_0 = p$ because by doing so we lose the capability of altering the support of the nominal measure.

In practice, a common problem of the relative entropy approach is that the denominator in Eq. 21 may not be integrable. To see this point, we examine the worst-case density function under the relative entropy approach:

$$q_{KL}^*(x) \propto p(x)e^{\theta V(x)} \quad (30)$$

The RHS of Eq. 30 may not be integrable if $V(x)$ increases too fast (or $p(x)$ decays too slowly as in the cases of heavy tails). As an example, we consider the worst-case variance problem where $V(x) = x^2$. If the nominal model follows an exponential distribution, then Eq. 30 is surely not integrable.

Using the proposed Wasserstein approach, however, the flexibility of choosing a proper prior q_0 helps us to bypass this issue. In fact, one may choose a prior distribution q_0 , different from the nominal distribution p , to guarantee that it decays sufficiently fast. According to Eq. 27, all we need to guarantee is that

$$q_0(y) \exp\left(\frac{V(y)}{\alpha} - \frac{c(x, y)}{\alpha\beta}\right) \quad (31)$$

is integrable w.r.t y . Fortunately, it is always possible to find some q_0 that satisfies this criteria. As a simple choice, we may set $q_0(y) \propto e^{-V(y)/\alpha}$ to ensure the integrability. Such choice makes Eq. 31 proportional to

$$\exp\left(-\frac{c(x, y)}{\alpha\beta}\right) \quad (32)$$

We suppose that the state space Ω is an Euclidean space with finite dimension and the transportation cost $c(x, y)$ is given by its Euclidean distance. Then for all $x \in \Omega$ Eq. 32 is integrable w.r.t y , for the integrand diminishes exponentially when y moves away from x .

In summary, formulating the problem using the relative entropy constraint Eq. 24 allows for flexibility of choosing a prior distribution q_0 . This is practically useful as one can avoid integrability issue by selecting a proper prior. This flexibility is not shared by the relative entropy approach as in Glasserman and Xu, which is regarded as a special case where the prior q_0 equals the nominal distribution p .

4 Application

4.1 Jump risk under diffusive nominal model

We start from a price process that takes the form of a geometric Brownian motion

$$dS_t = \mu S_t dt + \sigma S_t dW_t \quad (33)$$

The logarithmic return at time T follows a normal distribution:

$$x := \ln \left(\frac{S_T}{S_0} \right) \sim \mathcal{N} \left(\left(\mu - \frac{\sigma^2}{2} \right) T, \sigma^2 T \right) \quad (34)$$

When the volatility reaches zero, the return becomes deterministic and the distribution density is

$$p(x) = \lim_{\sigma \rightarrow 0} \frac{1}{\sqrt{2\pi T \sigma}} e^{-\frac{[x - (\mu - \sigma^2/2)T]^2}{2\sigma^2 T}} = \delta(x - \mu T) \quad (35)$$

In the case, model risk cannot be quantified using f -divergence. In fact, the nominal measure is a Dirac measure therefore no equivalent alternative measure exists. Under the KL divergence in particular, the worst-case measure is calculated by

$$\frac{p(x)e^{\theta V(x)}}{\int_{\Omega} p(x)e^{\theta V(x)} dx} = \delta(x - \mu T) \quad (36)$$

which is the same as the nominal measure. This is consistent with the Girsanov theorem for diffusion processes which states that the drift term is altered by some amount proportional to the volatility, i.e. $\tilde{\mu} = \mu - \lambda\sigma$. When the volatility under the nominal model decreases to zero, the alternative measure becomes identical to the nominal measure.

Approaches based on f -divergence excludes the existence of model risk given a zero volatility. This is, however, not true in practice, as the nominal diffusion process may still “regime-switch” to some discontinuous process. In fact, to quantify risks, one usually take into account the possibility of discontinuous changes of state variables (i.e. “jumps”). Using the Wasserstein approach, quantifying such jump risk becomes possible, even if the nominal model is based on a pure diffusion process. Substituting Eq. 35 into Eq. 21 gives the worst-case distribution (see Appendix for details)

$$q_W(x) = \frac{\exp \left(\frac{V(x)}{\alpha} - \frac{c(x, \mu T)}{\alpha\beta} \right)}{\int_{\Omega} \exp \left(\frac{V(y)}{\alpha} - \frac{c(y, \mu T)}{\alpha\beta} \right) dy} \quad (37)$$

Notice that Eq. 37 is suitable for any application where the nominal model is given by a Dirac measure. Under f -divergence, the limitation to equivalent measures keeps the nominal model unchanged. The Wasserstein approach, on the other hand, relaxes such limitation, allowing for a worst-case model that differs from a Dirac measure. This allows us to measure risk in deterministic variable or process. A particularly interesting example is the quadratic variation process, which is deemed as deterministic under the Black-Scholes model. We will discuss this in detail later with regards to the model risk in dynamic hedging.

To illustrate Eq. 37, we consider the expected value of x under the worst-case scenario. This problem is formulated using Eq. 6 with a linear loss function $V(x) = x$. We further assume a quadratic transportation cost function $c(x, y) = (x - y)^2$. The worst-case

distribution given by Eq. 37 turns out to be

$$q_W(x) = \frac{1}{\sqrt{\pi\alpha\beta}} e^{-\frac{(x-\mu T-\beta/2)^2}{\alpha\beta}} \quad (38)$$

One can see that the worst-case scenario is associated with a constant shift of the mean (by $-\beta/2$), even if the nominal measure is deterministic (i.e. Dirac). The change in mean is also associated with a proportional variance (i.e. $\alpha\beta/2$), if α is assigned a positive value. The resulted normal distribution, with a finite variance, is a reflection of model ambiguity. This is in contrast with approaches based on f -divergences, which are incapable of altering the nominal model as its support includes only a single point.

4.2 Volatility Risk and Variance Risk

In this section, we consider the risk of volatility uncertainty given the nominal Black-Scholes model. When an option approaches maturity, the nominal measure (on the price of its underlying asset) becomes close to a Dirac measure. This is visualised by the normal distribution of return narrowing in a rate of \sqrt{t} . When the time to maturity $t \rightarrow 0$, the normal distribution shifts to a Dirac distribution with zero variance.

Under the Kullback-Leibler divergence (or any f -divergence), any model risk vanishes when the nominal model converges to a Dirac measure. As a result, on a short time to maturity a sufficient amount of variance uncertainty can only be produced with a large cost (parametrised by θ). To illustrate this point, consider a normal distribution (say Eq. 35 before taking the limit). For the purpose of measuring the variance risk, we need to adopt a quadratic loss function $V(x) = x^2$. Under the Kullback-Leibler divergence, the variance of the worst-case distribution is given by[4]

$$\sigma_{KL}^2 T = \frac{\sigma^2 T}{1 - 2\theta\sigma^2 T} \quad (39)$$

When time to maturity $T \rightarrow 0$, the worst-case volatility $\sigma_{KL} \rightarrow \sigma$ with a fixed θ . This is not consistent with what we saw on the market. In fact, on short time to maturity the fear of jumps can play an important role. Such fear of risks is priced into options and variance swaps termed as the volatility (or variance) risk premium.




Fig 2: Worst-case volatility as a function of time under the (a) Wasserstein approach, (b) KL divergence.

The volatility (or variance) risk premium can be considered as the compensation paid to option sellers for bearing the volatility risks[8, 9]. It is practically quantified as the difference between the implied volatility (or variance) and the realised volatility (or variance). As it is priced based on the volatility risk, its quantity is directly linked to the risk associated with the nominal measure used to model the underlying asset. Therefore by analyzing the term structure of such premium, one can get some insight into the worst-case volatility risk. Under the assumption of diffusive price dynamics, Carr and Wu has developed a formula for the at-the-money implied variance[10]. Illustrated in Fig. 2, the formula matches well with the empirical data[11] for maturities longer than 3 months. For maturities shorter than 3 months, however, the formula seems to underestimate the variance risk premium. Other empirical work also shows that option buyers consistently pay higher risk premium for shorter maturity options[9].

The underestimation of volatility risk premium on short maturity is an intrinsic problem with diffusive models. Indeed, the work mentioned above reveals the importance of quantifying jumps on short time to maturity. Other work shows that the risk premium due to jumps is fairly constant across different maturities[12]. This implies a very different time dependency from that due to continuous price moves (Eq. 39). In fact, any approach based on f -divergence is incapable of producing sufficient model risk on $t \rightarrow 0$, suggesting a decaying term structure of risk premium. On the other hand, the Wasserstein approach does not suffer from this issue. In fact, it produces a worst-case volatility that has little time dependence (Fig. 2). Therefore, the Wasserstein approach provides a particularly useful tool for managing the variance risk and quantifying its risk premium on short time to maturity.

With the Wasserstein approach, the worst-case variance takes the form of (see Appendix B)

$$\sigma_W^2 T = \frac{\sigma^2 T}{(1 - \beta)^2} + \frac{\alpha \beta}{2(1 - \beta)} \quad (40)$$

The Wasserstein approach provides a worst-case variance that is independent of the time to maturity. It scales the nominal variance by a constant factor $(1 - \beta)^{-2}$. In addition, it introduces a constant extra variance $\alpha \beta / (1 - \beta)$. The extra variance term is modulated by

the parameter α . If we set α to zero, then the worst-case volatility σ_W is merely a constant amplification of the nominal volatility σ . This model risk measure, however, may not be sufficient if the nominal volatility is very close to zero. The extra variance term serves to account for the extra risks (e.g. jumps) that are not captured by the nominal volatility.




Fig 3: Worst-case volatility as a function of time under the (a) Wasserstein approach, (b) KL divergence.




Fig 4: (a) Percentage volatility premium as the surplus of worst-case volatility over the nominal volatility and (b) Worst-case distributions under the two approaches.

4.3 Model Risk in Portfolio Variance

The Wasserstein approach can be applied to quantify the risk associated with modelling the variance of a portfolio, assuming the asset returns follow a multivariate normal distribution. Suppose there are n assets under consideration and their returns are reflected by a state vector x : $x \in \mathcal{V}$ where \mathcal{V} is a n -dimensional vector space. For generality, we consider the following target function $V : \mathcal{V} \rightarrow \mathbb{R}_+$

$$V(x) = x^T A x \quad (41)$$

where A is a positive-definite symmetric matrix. If we replace x by $x' = x - E(x)$ and A by ww^T , then the expected value of the target function reflects the portfolio variance:

$$E[V(x)] = E(x^T w w^T x) = w^T \Sigma w \quad (42)$$

where w is the vector of compositions in the portfolio. Σ is the covariance matrix of the normally distributed asset returns (under the nominal model).

To find the worst-case model using the Wasserstein approach, we need to first define a metric in the vector space \mathcal{V} . Suppose the vector space is equipped by a norm $\|x\|$ then the metric is naturally defined by $c(x, y) = \|x - y\|$. Here we focus on the kind of norm that has an inner-product structure:

$$\|x\| = \sqrt{x^T B x}, \quad \forall x \in \mathcal{V} \quad (43)$$

where B is a positive-definite symmetric matrix (constant metric tensor). The resulted worst-case distribution is still multivariate normal, with the vector of means and covariance matrix replaced by (see Appendix D for derivation)

$$\mu_W = (B - \beta A)^{-1} B \mu \quad (44)$$

$$\Sigma_W = (B - \beta A)^{-1} B \Sigma B (B - \beta A)^{-1} + \frac{\alpha \beta}{2} (B - \beta A)^{-1} \quad (45)$$

Apart from a constant term that vanishes if assigning zero to the parameter α , the worst-case distribution is transformed from the nominal distribution via a measure-preserving linear map (see Appendix D). This result is more intuitive than the result obtained using the KL divergence, given by[4]

$$\mu_{KL} = (I - 2\theta \Sigma A)^{-1} \mu \quad (46)$$

$$\Sigma_{KL} = (I - 2\theta \Sigma A)^{-1} \Sigma \quad (47)$$

Fig. 5 provides an example showing that the worst-case distribution is indeed a measure-preserving transform with the Wasserstein approach.




Fig 5: Multivariate nominal distributions (a) nominal model, (b) worst case under the KL divergence, (c) worst case under the Wasserstein approach (as a measure-preserving transform).

The constant term reflects residual uncertainty when the nominal model has vanishing variances. This term is especially useful when some of the assets are perfectly correlated (either 1 or -1) and the vector space \mathcal{V} is not fully supported by the nominal measure. In this case, the Wasserstein approach provides results that differ significantly from the f -divergence approach. In particular, approaches based on KL divergence (or any f -divergences) cannot alter the support, they merely reweight the states within the support. This is illustrated in Fig. 6, where two assets are perfectly correlated. The nominal model shown in (a) provides a measure supported by a one-dimensional vector subspace of \mathcal{V} . The worst-case measure under the KL divergence is supported by the same subspace, as




Fig 6: Multivariate nominal distributions (a) nominal model, (b) worst case under the KL divergence, when the support is a low-dimensional subspace. Worst-case multivariate nominal distributions under the Wasserstein approach (c) $\theta = 0$ (d) $\theta = 0.5$.

illustrated in (b). This conclusion can actually be derived from the worst-case measure given by Eq. 47 (see Appendix E for proof).

On the other hand, the Wasserstein approach is capable of examining measures supported by other vector subspaces. We first ignore the constant variance term by setting α to zero in Eq. 45. The Wasserstein approach “rotates” the original support by applying linear maps to the nominal measure. In the case illustrated by Fig. 6(c), essentially all measures supported by a one-dimensional vector subspace are within the scope of the approach (see Appendix F for proof). Among those measures the Wasserstein approach picks the worst one, supported by a vector subspace different from the original one. It essentially searches for the optimal transform over the entire space. In practice, we may want to account for the risk associated with the assumption of perfect correlation. This is accomplished by assigning positive value to α , allowing the distribution to “diffuse” into the entire vector space as illustrated in Fig. 6(d).

It is worthwhile to notice that the Wasserstein approach also has a practical advantage over the approach based on KL divergence. If we examine the worst-case variances resulted from the two approaches, Eq. 45 and 47, we can find that their positive definiteness is not guaranteed. This requires practitioners to carefully parametrise either approach to ensure the positive definiteness. However, under KL divergence the positive definiteness is dependent on the original covariance matrix. This makes it harder to parametrise and generalise the approach. In cases where the asset returns has time-varying correlations, one may need to switch parameters (θ) to ensure a positive definite matrix. On the other hand, the Wasserstein approach only requires $B - \beta A$ to be positive-definite, independent of the covariance matrix Σ . The nominal model thus no longer affects its feasibility.

4.4 Robust Portfolio Optimisation and Correlation Risk

In modern portfolio theory, one consider n securities with the excess logarithmic returns being multivariate normally distribution, i.e. $X \sim \mathcal{N}(\mu, \Sigma)$. The standard mean-variance optimisation is formulated by

$$\min_a a^T \Sigma a \quad (48)$$

$$s.t. \mu^T a = C \quad (49)$$

where $a \in \mathbb{R}^n$ is the vector of portfolio weights. It can take any values assuming it is always possible to borrow or lend at the risk-free rate. The problem is solved by introducing a Lagrange multiplier λ :

$$a^* = \frac{\lambda}{2} \Sigma^{-1} \mu \quad (50)$$

The optimal portfolio weight a^* depends on λ . However, the Sharpe ratio of the optimal portfolio is independent of λ :

$$\frac{a^{*T} \mu}{\sqrt{a^{*T} \Sigma a^*}} = \sqrt{\mu^T \Sigma^{-1} \mu} \quad (51)$$

The nominal model assumes a multivariate normal distribution $\mathcal{N}(\mu, \Sigma)$. The worst-case model is an alternative measure dependent on the security positions a . To formulate the problem of worst-case measure, we may first express the mean-variance optimisation problem by

$$\min_a E [(x - \mu)^T a a^T (x - \mu) - \lambda x^T a] \quad (52)$$

where the expectation is taken under the nominal measure. Taking into account the model risk, we may formulate a robust version of Eq. 52 that is consistent with literature work [4]:

$$\min_a \max_{Q \in \mathcal{M}} E^Q [(X - \mu)^T a a^T (X - \mu) - \lambda X^T a] \quad (53)$$

where \mathcal{M} is the space of alternative measures constrained by different criteria. For the approached based on the Kullback-Leibler divergence, the constraint is given by a maximum amount of relative entropy w.r.t the nominal model (i.e. relative entropy budget). Under the Wasserstein approach, the constraints are given by Eq. 12 and 16.

To solve the inner problem of Eq. 53, we may further simplify the problem to

$$\begin{aligned} & \max_{Q \in \mathcal{M}} E^Q [(X - \mu)^T a a^T (X - \mu) - \lambda X^T a] \\ &= \max_{Q \in \mathcal{M}} E^Q [(X - \mu - k)^T a a^T (X - \mu - k)] - \lambda \mu^T a - \frac{\lambda}{4} \end{aligned} \quad (54)$$

where k is a vector that satisfies $a^T k = \lambda/2$. It is noted that the problem is at most an approximate as the change of measure would also alter the mean from μ to μ' . The variance should be calculated by $E^{Q(m)} [(X - \mu')^T a a^T (X - \mu')]$. However, the difference is proportional to $(\mu' - \mu)^2$ and is thus secondary. The solution to Eq. 53 is also multivariate under both KL divergence (see Appendix G) and under the Wasserstein metric (see Appendix H). The two approaches result in robust MVO portfolios with different weights

(up to the first order w.r.t θ or β):

$$\begin{aligned} a_{KL}^* &= \left(\frac{\lambda}{2} - \frac{\theta\lambda^3}{2} (1 + \mu^T \Sigma^{-1} \mu) \right) \Sigma^{-1} \mu \\ a_W^* &= \left(\frac{\lambda}{2} - \frac{\beta\lambda^3}{4} \mu^T \Sigma^{-1} B^{-1} \Sigma^{-1} \mu - \frac{\beta\lambda^3}{4} (1 + \mu^T \Sigma^{-1} \mu) \Sigma^{-1} B^{-1} \right) \Sigma^{-1} \mu \end{aligned} \quad (55)$$

Comparing Eq. 55 with the normal MVO portfolio given by Eq. 50, we can see that the robust MVO portfolios provide first-order corrections, resulting in more conservative asset allocation in general.

Despite of being more conservative, a_{KL}^* is in fact parallel to the normal MVO portfolio a^* . As a result, the robust MVO portfolio does not change the relative weights of component assets. In fact, all the weights are reduced by the same proportion ($c < 1$) to account for model risk. This is, however, inappropriately account for the correlation risk. For example, two highly-correlated assets have extremely high weights in the nominal MVO portfolio. Because of the correlation risk, we would expect the robust MVO portfolio to assign them lower weights relative to other assets. This is, however, no longer an issue for a_W^* . In fact, a_W^* not only reduces the overall portfolio weights in order to be more conservative, but also adjust the relative weights of component assets for a less extreme allocation. One may notice that the term inside the bracket of the expression for a_W^* is a square matrix (see Eq. 55), which serves to linearly transform the vector of portfolio weights. By adjusting their relative weights, Eq. 143 correctly accounts for the correlation risk (see Appendix H for details).

The robust optimal portfolios parametrised by λ allows us to plot the robust capital allocation line (CAL). Unlike the standard CAL, it is no longer a straight line and the Sharpe ratio is now dependent on λ .

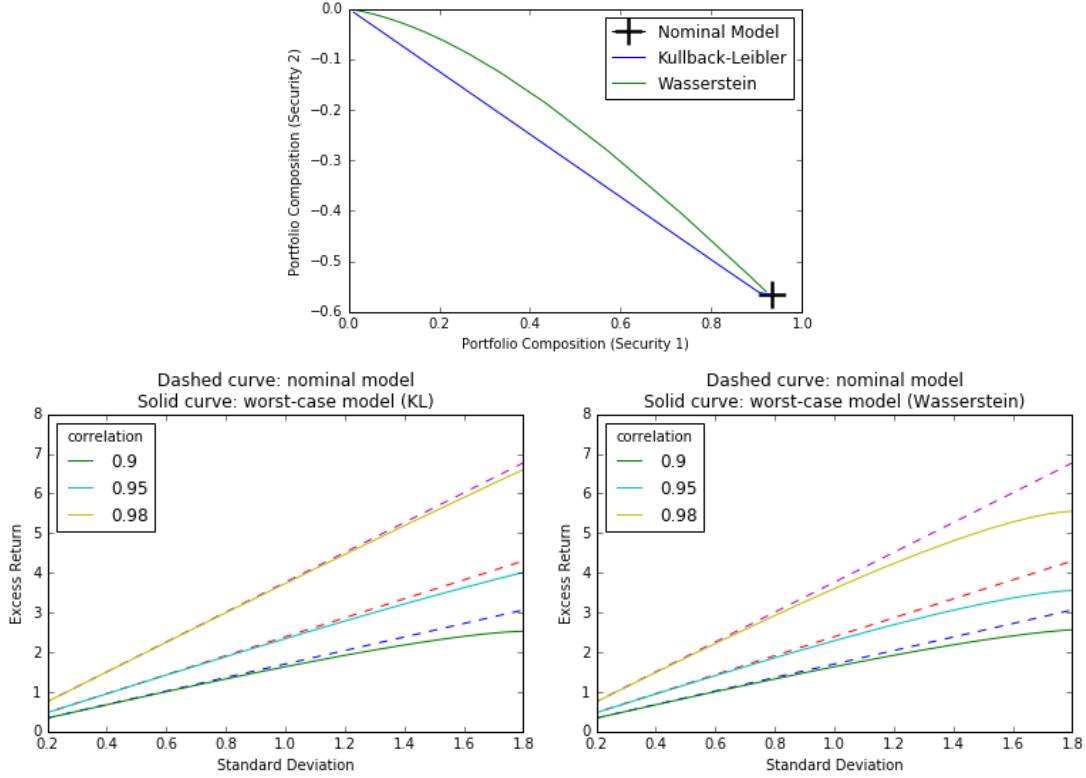


Fig 7: The normalised optimal composition of a portfolio consisting of two securities, calculated by a^* divided by $\lambda/2$. The normalised optimal composition under the nominal model is given by a constant vector $\Sigma^{-1}\mu$, while those under the worst-case models are dependent on λ . In particular, the Kullback-Leibler approach reduces both compositions proportionally, while the Wasserstein approach reduces compositions in a nonlinear way.

Under the nominal model, the optimal composition of a portfolio is given by $\lambda\Sigma^{-1}\mu/2$. The proportionality of this solution suggests that we should double the weights if the expected excess return doubles. However, this may end up with excessive risk due to an increase in leverage. Model risk is the major source of risks here, as we are unsure if the expected excess return and the covariance matrix correctly reflect the return statistics in the future (for a given holding period). Since higher leverage implies more severe model risk, increasing leverage proportionally is in fact sub-optimal under the worst-case model.

Eq. 55, on the other hand, provides the optimal solutions under the respective model risk approaches. The robustness of these solutions allows practitioners to allocate assets in a safer way. It is shown in Fig. 7 that the normalised optimal compositions reduce with λ . This is because a larger λ indicates higher leverage, and hence the optimal composition is reduced further away from that of the nominal model. The normalised optimal compositions approach zero as λ increases. In Fig. 7, the compositions of both securities get reduced proportionally under the KL approach. Using the Wasserstein approach, on the other hand, allows the compositions to move in a non-linear way.

In this example, we have two highly correlated ($\rho = 0.5$) stocks but with very different expected excess returns (Stock 1 0.65 and Stock 2 -0.1). Because of the high correlation, we can profit from taking the spread (long Stock 1 and short Stock 2). Under the nominal model, taking a spread of a highly correlated pair does not add too much risk. However, the true risk could be underestimated due to the existence of model risk. The spread is

more sensitive to model risk than an overall long position, thus requires reduction when optimising with model risk. This point is well reflected by the non-linearity of the optimal compositions under the Wasserstein approach. We reduces the position of the spread more than the long position of Stock 1 (or the overall long position). In the KL approach, however, we reduces the spread position and the overall long position at the same pace.

The effect of robust optimality under the worst-case model is most significant when the nominal model is close to have a low-dimensional support. A low-dimensional support means that the covariance matrix does not have the full rank. Put it in a practical way, there exists a risk-free portfolio with non-zero compositions in risky assets. In this case, there is arbitrage opportunity that has close-to-zero risk but high excess returns. The optimal portfolio under the nominal model could be unrealistically optimistic.

Fig. 7(b) and (c) illustrate an example of two securities with a high correlation. Under the nominal model, the Sharpe ratio (slope of the excess return vs std dev lines) increases quickly with the correlation coefficient. This results from taking excessive positions in the spread (long the one with higher Sharpe ratio and short the other). From Fig. ??, it is shown that the KL approach cannot solve this issue systematically. In fact, when the correlation increases, the capital allocation line under the worst-case model is even closer to the nominal one. On the other hand, the Wasserstein approach does provide the correct adjustment. The robust capital allocation line given by the Wasserstein approach deviates more from the nominal straight line on an increasing correlation.

This difference is a direct result in their capabilities of altering the support of the nominal measure. The KL approach cannot alter the support. So an arbitrage relation under the nominal measure may persist under the worst-case measure. On the other hand, the Wasserstein search for a support for the worst-case measure. It breaks the arbitrage opportunity by transforming the support to a different vector subspace.

4.5 Model Risk in Dynamic Hedging

The hedging error is measured by the absolute profit-and-loss (PnL) of a dynamically hedged option until its maturity. Using the Black-Scholes model as the nominal model, the hedging risk decreases with the hedging frequency. Ideally if hedging is done continuously, then the hedging error is zero almost surely. This is true even under alternative measures, as long as they are equivalent to the nominal model. The underlying reason is that the quadratic variation does not change under all equivalent measures. In fact, if we consider a geometric Brownian motion:

$$dS_t = \mu S_t dt + \sigma S_t dW_t \quad (56)$$

The quadratic variation $[\ln S]_t = \int_0^t \sigma_s^2 ds$ almost surely. Therefore the equation holds under all equivalent measures. Given the Black-Scholes price of an option $C_t = C(t, S_t)$, the PnL of a continuously hedged portfolio between time 0 and T is

$$\begin{aligned} & \int_0^T dC_t - \int_0^T \frac{\partial C_t}{\partial S_t} dS_t \\ &= \int_0^T \left(\frac{\partial C_t}{\partial t} dt + \frac{S_t^2}{2} \frac{\partial^2 C_t}{\partial S_t^2} d[\ln S]_t \right) = 0 \end{aligned} \quad (57)$$

where the last equality results from the Black-Scholes partial differential equation.

Since any f -divergence is only capable of searching over equivalent alternative measures, the worst-case hedging error given by these approaches has to be zero on continuous

hedging frequency. One can imagine that as hedging frequency increases, the worst-case hedging risk decreases towards zero (Fig. 8(b)). This is, however, inconsistent with practitioners' demand for risk management. In fact, if the volatility of the underlying asset differs from the nominal volatility, then Eq. 57 no longer holds. Such volatility uncertainty is a major source of hedging risk, and thus has to be measured and managed properly. The most straightforward way of doing that is to assume a distribution of volatility, and then run a Monte Carlo simulation to quantify the hedging error (Fig. 8(a)).




Fig 8: (a) Worst-case hedging risk under the KL divergence, and (b) hedging risk simulated by randomly sampling volatilities.

Despite of its simplicity, volatility sampling is a parametric approach, for it is only capable of generating alternative Black-Scholes models with different parameter. This approach cannot account for alternatives such as local volatility models or stochastic volatility models. This calls for a non-parametric approach relying on the formulation given in Eq. 6.

We have already seen that using approaches based on f -divergence one cannot correctly quantify the hedging risk. The Wasserstein approach, on the other hand, does not have this issue, for it is capable of searching over non-equivalent measures. Using Monte Carlo simulation, we obtain the worst-case hedging risk under the Wasserstein approach (see Fig. 9). Compared to the approach based on Kullback-Leibler divergence (Fig. 8(b)), the hedging risk given by the Wasserstein approach is more consistent with the simulated results using volatility sampling (Fig. 8(a)). In the limit of continuous hedging, the Wasserstein approach results in a worst-case risk slightly higher than volatility sampling, for it may involve jumps that cannot be hedged.

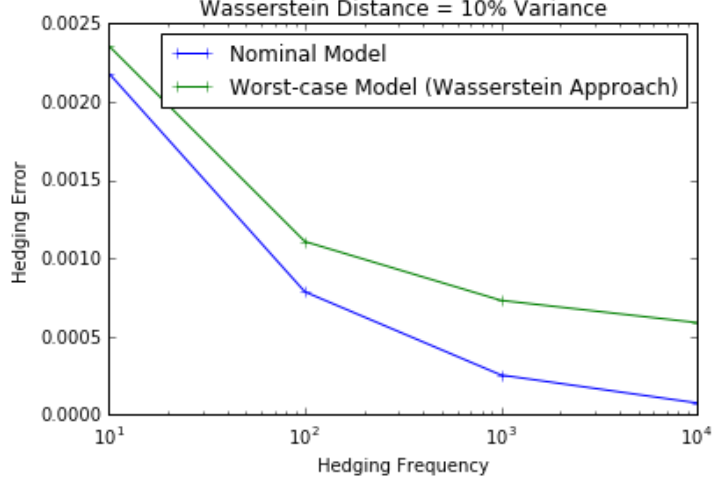


Fig 9: (a) Worst-case hedging risk under the Wasserstein approach.

In practice, the Wasserstein approach requires some tricks as fully sampling the infinite-dimensional path space is impossible. Therefore only paths close to the sampled paths (under the nominal measure) are sampled, as the importance of an alternative path decays exponentially with its distance to these sampled paths. This point is shown in Fig. 10(a), in which the alternative paths are illustrated by the crosses close to the nominal sampled paths (dots). By increasing the average distance of the alternative paths to the nominal paths, the hedging risk is increased until convergence (Fig. 10(b)).

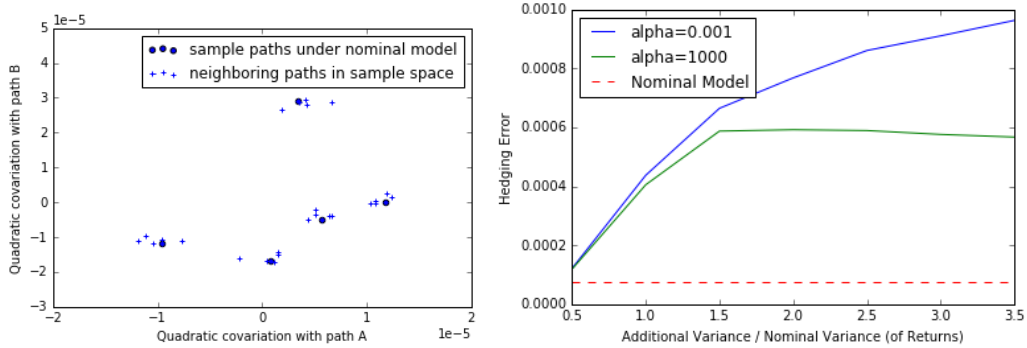


Fig 10: (a) Sample paths generated for the Wasserstein approach, (b) convergence of the worst-case hedging risk.

Here we list the procedure of the Monte Carlo simulation described in the last paragraph:

1. create N sample paths from the nominal model
2. For each sample paths, creat M sample paths by deviating X_t by a normally distributed random variable $\mathcal{N}(0, \sigma^2)$
3. collect all MN sample paths and the original N paths, we have $N(M + 1)$ points in the path space. Calculate the hedging error for each of the $N(M + 1)$ paths.
4. Apply Eq. 21 to find the worst-case probability of each path where $d(X, Y) = [X - Y]$.
5. To find the (worst-case) hedging risk, we average the hedging errors of all $N(M + 1)$ paths, weighted by their worst-case probabilities.

6. Repeat steps 2-5 with a larger σ^2 . Continue to increase the deviation until the calculated hedging risk gets converged.

5 Conclusion

Non-parametric approaches to model risk measurement are theoretically sound and practically feasible. Adopting the Wasserstein distance allows us to further extend the range of legitimate measures from merely the absolutely continuous ones. This Wasserstein approach roots from the transportation theory and is well suited for the adversary interpretation of the model risk. In particular, it specifies the economic reality of the fictitious adversary with the capacity of parametrising the actual market structure. The Wasserstein approach may result in the worst-case model that is more robust, in the sense that it is no longer restricted by the support of the nominal measure. This is especially useful when the nominal measure is supported only by a subspace (for instance the volatility of a diffusion process or the prices of perfectly correlated assets). This approach has additional practical advantage due to its ability of guaranteeing integrability.

To further illustrate the Wasserstein approach, we propose four applications ranging from the single-asset variance risk and hedging risk to the multi-asset allocation problem. All the applications are connected in the sense that their nominal measures are (or close to be) supported by merely a subspace. In the example of single-asset variance risk, we look into the limit of small variance, i.e. when the time to maturity is close to zero (or the volatility close to zero). The Wasserstein approach is capable of jumping out of the family of diffusion processes, and account for the possibility of jumps. In the application of portfolio variance risk, the Wasserstein approach provides us with worst-case measure induced by a linear map, thus altering the support. Its advantage of dealing with multi-asset problems is even more apparent when treating with the asset allocation problem, in which the Wasserstein approach correctly accounts for the correlation risk. This approach results in a robust mean-variance optimal portfolio that adjusts the relative weights of the assets according to their correlations. It produces a curved capital allocation line, with the Sharpe ratio reduced by a larger amount on a higher standard deviation or a higher asset correlation. The final application is related to the hedging risk of a vanilla option. f -divergence is incapable of quantifying the risk associated with a continuously hedged position because its profit-and-loss is zero almost surely. The Wasserstein approach, on the other hand, leads to a positive hedging error in consistency with the simulation result. In conclusion, the Wasserstein approach provides a useful tool to practitioners who aim to manage risks and optimize positions accounting for model ambiguity.

6 Appendix

6.1 A. Derivation of Eq. 15

In this part, we derive the solution Eq. 15 to the problem expressed by Eq. 10-12. For simplicity, we denote the transition density $p_{Y|X}(y|x)$ by $\gamma_x(y) := \gamma(x, y)/p(x)$. This

transforms the problem into

$$\begin{aligned} & \sup_{\gamma_x \in \Gamma} \int_{\Omega} p(x) \left[\int_{\Omega} \gamma_x(y) V(y) dy \right] dx & (58) \\ & \text{s.t.} \int_{\Omega} p(x) \left[\int_{\Omega} \gamma_x(y) c(x, y) dy \right] dx \leq \eta \end{aligned}$$

where Γ is the space of probability density functions. The Karush-Kuhn-Tucker (KKT) condition in convex optimisation ensures the existence of a KKT multiplier λ such that the solution to Eq. 58 also solves

$$\sup_{\gamma_x \in \Gamma} \int_{\Omega} p(x) \left\{ \int_{\Omega} \gamma_x(y) [V(y) - \lambda c(x, y)] dy \right\} dx \quad (59)$$

The solution to Eq. 59 is a δ -function transition density $\gamma_x^*(y) = \delta(y - y^*(x))$, resulting in a transportation plan

$$\gamma^*(x, y) = p(x) \delta(y - y^*(x)) \quad (60)$$

where

$$y^*(x) = \arg \max_{y \in \Omega} [V(y) - \lambda c(x, y)] \quad (61)$$

The solution to the model risk problem is expressed either by a transportation plan (Eq. 60) or a transportation map (Eq. 61). It is noted that $\lambda = 0$ is a trivial case that we will not consider. To be consistent with the main result Eq. 21, we replace λ by its inverse $\beta = \lambda^{-1}$:

$$y^*(x) = \arg \max_{y \in \Omega} \left[V(y) - \frac{c(x, y)}{\beta} \right] \quad (62)$$

6.2 B. Derivation of Eq. 21 and 25

Eq. 21 is the solution of the problem formulated by Eq. 10-12 plus the additional entropy constraint Eq. 16. As in Appendix A, we introduce KKT multipliers λ and α . This converts the original constrained superimum problem to the following dual problem (same as in Appendix A we denote the transition density by $\gamma_x(y)$):

$$\begin{aligned} & \inf_{\beta, \theta \in \mathbb{R}^+} \sup_{\gamma} \int_{\Omega \times \Omega} \gamma(x, y) (V(y) - \lambda [c(x, y) - \eta] - \alpha [\ln \gamma(x, y) - \mu]) dx dy & (63) \\ & = \inf_{\beta, \theta \in \mathbb{R}^+} \left(\int_{\Omega} p(x) dx \left[\sup_{\gamma_x} \int_{\Omega} \gamma_x(y) (V(y) - \lambda c(x, y) - \alpha \ln \gamma_x(y)) dy \right] \right. \\ & \quad \left. + \lambda \eta + \alpha \left[\mu - \int_{\Omega} \ln p(x) dx \right] \right) \end{aligned}$$

Same as the relative entropy approach proposed by Glasserman and Xu, we derive a closed-form solution to the inner part of the problem:

$$\sup_{\gamma_x} \int_{\Omega} \gamma_x(y) (V(y) - \lambda c(x, y) - \alpha \ln \gamma_x(y)) dy \quad (64)$$

It is noted that Eq. 64 asks for the supermium w.r.t the density function p_x for a given $x \in \Omega$. The solution to this problem is given by (for consistency we replace λ by its

inverse γ):

$$\gamma_x^*(y) = \frac{\exp\left(\frac{V(y)}{\alpha} - \frac{c(x,y)}{\alpha\beta}\right)}{\int_{\Omega} \exp\left(\frac{V(y)}{\alpha} - \frac{c(x,y)}{\alpha\beta}\right) dy} \quad (65)$$

The worst-case probability density function is the marginal distribution of y , induced by the transition density function $\gamma_x^*(y)$:

$$\begin{aligned} p^*(y) &= \int_{\Omega} p(x) \gamma_x^*(y) dx \\ &= \int_{\Omega} dx \frac{p(x) \exp\left(\frac{V(y)}{\alpha} - \frac{c(x,y)}{\alpha\beta}\right)}{\int_{\Omega} \exp\left(\frac{V(y)}{\alpha} - \frac{c(x,y)}{\alpha\beta}\right) dy} \end{aligned} \quad (66)$$

Eq. 25 is derived in a similar way. Since we lift the entropy constraint Eq. 16 into a relative entropy constraint Eq. 19, the inner problem Eq. 64 requires slight modification:

$$\sup_{\gamma_x} \int_{\Omega} \gamma_x(y) \left(V(y) - \lambda c(x,y) - \alpha \ln \frac{\gamma_x(y)}{q_0(y)} \right) dy \quad (67)$$

This problem has the same formulation as the supermium problem given in Glasserman and Xu's work, and therefore shares the same solution

$$\gamma_x^*(y) = \frac{q_0(y) \exp\left(\frac{V(y)}{\alpha} - \frac{c(x,y)}{\alpha\beta}\right)}{\int_{\Omega} q_0(y) \exp\left(\frac{V(y)}{\alpha} - \frac{c(x,y)}{\alpha\beta}\right) dy} \quad (68)$$

This equation differs from Eq. 65 merely by a prior distribution q_0 . It takes Eq. 65 as its special case where q_0 is a uniform distribution. Marginalizing the transition density Eq. 68 gives the worst-case distribution shown in Eq. 25.

6.3 C. Jump Risk and Variance Risk

Under a diffusive model, the logarithmic return of an asset follows a normal distribution with mean of μT and variance of $\sigma^2 T$, where μ is the drift coefficient, σ is the volatility and T is the time to maturity. The probability density function of the return x is

$$p(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\left(\frac{(x-\mu T)^2}{2\sigma^2 T}\right)^2} \quad (69)$$

Applying Eq. 21, one may obtain the probability density function of the worst-case measure, assuming a linear loss function $V(x) = x$ and a quadratic transportation cost function $c(x,y) = (x-y)^2$,

$$\begin{aligned} q^*(y) &\propto \int_{\Omega} p(x) \exp\left(\frac{y-x}{\alpha} - \frac{(x-y)^2}{\alpha\beta}\right) dx \\ &= \int_{\Omega} \exp\left(\frac{y-x}{\alpha} - \frac{(x-y)^2}{\alpha\beta} - \frac{(x-\mu T)^2}{2\sigma^2 T}\right) dx \\ &= \exp\left(-\frac{(y-\mu T - \beta/2)^2}{2\sigma^2 T + \alpha\beta}\right) \end{aligned} \quad (70)$$

Unlike the result given by the KL divergence, Eq. 70 not only shift the mean the distribution but also enlarge the variance as a result of additional uncertainty. On $\sigma \rightarrow 0$, the worst-case measure is no longer a Dirac measure, showing consideration of jump risks:

$$\lim_{\sigma \rightarrow 0} q^*(y) \propto \exp\left(-\frac{(x - \mu T - \beta/2)^2}{\alpha\beta}\right) \quad (71)$$

This gives Eq. 38. Alternatively, one may first derive Eq. 37 followed by substituting $V(x) = x$ to get Eq. 38. Eq. 37 is derived by substituting $p(x) = \delta(x - \mu T)$ into Eq. 21:

$$q^*(y) = \int_{\Omega} \delta(x - \mu T) \frac{\exp\left(\frac{V(y)}{\alpha} - \frac{(x-y)^2}{\alpha\beta}\right)}{\int_{\Omega} \exp\left(\frac{V(y)}{\alpha} - \frac{(x-y)^2}{\alpha\beta}\right) dy} dx \quad (72)$$

$$= \frac{\exp\left(\frac{V(y)}{\alpha} - \frac{(y-\mu T)^2}{\alpha\beta}\right)}{\int_{\Omega} \exp\left(\frac{V(y)}{\alpha} - \frac{(y-\mu T)^2}{\alpha\beta}\right) dy} \quad (73)$$

Now we adopt a quadratic type of loss function, $V(x) = (x - \mu T)^2$, following a procedure similar to Eq. 70 we get

$$q^*(y) \propto \exp\left(-\frac{(y - \mu T)^2}{\frac{2\sigma^2 T}{(1-\beta)^2} + \frac{\alpha\beta}{(1-\beta)}}\right) \quad (74)$$

the variance of the worst-case measure is

$$\sigma_W^2 T = \frac{\sigma^2 T}{(1-\beta)^2} + \frac{\alpha\beta}{2(1-\beta)} \quad (75)$$

as provided in Eq. 40. We may verify that the measure Q^* given by Eq. 74 does provide the largest variance among all the legitimate alternative measures. In fact, the variance of x under Q^* is

$$E^{Q^*} \left[\left(x - E^{Q^*}(x) \right)^2 \right] = E^{Q^*} \left[(x - \mu T)^2 \right] \quad (76)$$

According to the definition of the worst-case model, for all $Q \in \mathcal{M}$ (the space of legitimate alternative measures) we have

$$E^{Q^*} \left[(x - \mu T)^2 \right] \geq E^Q \left[(x - \mu T)^2 \right] \quad (77)$$

$$= E^Q \left[(x - E^Q(x))^2 \right] + (E^Q(x) - \mu T)^2 \quad (78)$$

$$\geq E^Q \left[(x - E^Q(x))^2 \right] \quad (79)$$

This confirms that Eq. 75 is indeed the worst-case (maximum) variance.

6.4 D. Worst-case Portfolio Variance

To find the portfolio variance under the worst-case scenario, we need to formulate the problem using Eq. 6 with a loss (target) function given by Eq. 41. The worst-case measure may be evaluated by substituting the loss function into Eq. 21. In this section we will show the calculation step by step. First, we need to specify the transport cost function $c(x, y)$ as

the inner product introduced in Eq. 43:

$$c(x, y) = \|y - x\|^2 = (y - x)^T B (y - x) \quad (80)$$

Then we evaluate the following part in Eq. 21:

$$\begin{aligned} & \exp\left(\frac{V(y)}{\alpha} - \frac{c(x, y)}{\alpha\beta}\right) \\ &= \exp\left(\frac{y^T A y}{\alpha} - \frac{(y - x)^T B (y - x)}{\alpha\beta}\right) \\ &= \exp\left(\frac{1}{\alpha\beta} x^T B (B - \beta A)^{-1} B x \right. \\ & \quad \left. - \frac{1}{\alpha\beta} (y - (B - \beta A)^{-1} B x)^T (B - \beta A) (y - (B - \beta A)^{-1} B x)\right) \end{aligned} \quad (81)$$

Remember that both A and B are symmetric, positive-definite matrices. Fixing x , Eq. 81 is proportional to the probability density function of a multivariate normal variable Y , with its mean and covariance matrix

$$E(Y) = (B - \beta A)^{-1} B x \quad (82)$$

$$\Sigma(Y) = \frac{\alpha\beta}{2} (B - \beta A)^{-1} \quad (83)$$

This means that after normalization w.r.t y , Eq. 81 gives exactly the probability density function of Y . We may write this down explicitly by noticing that y lives in the n -dimensional vector space, i.e. $\Omega = \mathcal{V}$:

$$\begin{aligned} & \frac{\exp\left(\frac{V(y)}{\alpha} - \frac{c(x, y)}{\alpha\beta}\right)}{\int_{\mathcal{V}} \exp\left(\frac{V(y)}{\alpha} - \frac{c(x, y)}{\alpha\beta}\right) dy} \\ &= (2\pi)^{-\frac{n}{2}} \sqrt{\frac{\alpha\beta}{2} |B - \beta A|} \exp\left(- (y - (B - \beta A)^{-1} B x)^T \frac{(B - \beta A)}{\alpha\beta} (y - (B - \beta A)^{-1} B x)\right) \end{aligned} \quad (84)$$

Now we need to evaluate the product of Eq. 84 and the nominal distribution $p(x)$. The nominal distribution is multivariate normal with mean μ and covariance matrix Σ :

$$p(x) = \frac{(2\pi)^{-\frac{n}{2}}}{|\Sigma|} \exp\left(-\frac{1}{2} (x - \mu)^T \Sigma^{-1} (x - \mu)\right) \quad (85)$$

The product contains many terms of x and y . One may re-arrange the terms to isolate quadratic and linear terms of x :

$$\begin{aligned}
& \frac{p(x) \exp\left(\frac{V(y)}{\alpha} - \frac{c(x,y)}{\alpha\beta}\right)}{\int_{\mathcal{Y}} \exp\left(\frac{V(y)}{\alpha} - \frac{c(x,y)}{\alpha\beta}\right) dy} \\
& \propto \exp\left(-\frac{1}{\alpha\beta} (y - (B - \beta A)^{-1} Bx)^T (B - \beta A) (y - (B - \beta A)^{-1} Bx) \right. \\
& \quad \left. - \frac{1}{2} (x - \mu)^T \Sigma^{-1} (x - \mu)\right) \\
& = \exp\left(-\frac{1}{\alpha\beta} [(x - K\mu - Ly)^T M (x - K\mu - Ly) - (K\mu + Ly)^T M (K\mu + Ly)] \right. \\
& \quad \left. - \frac{1}{\alpha\beta} y^T (B - \beta A) y\right) \tag{86}
\end{aligned}$$

where

$$\begin{aligned}
M & := B(B - \beta A)^{-1} B + \frac{\alpha\beta}{2} \Sigma^{-1} \\
K & := \frac{\alpha\beta}{2} M^{-1} \Sigma^{-1} \\
L & := M^{-1} B
\end{aligned} \tag{87}$$

Fixing y , Eq. 86 is proportional to the probability density function of a multivariate normal variable X where

$$E(X) = K\mu + Ly \tag{88}$$

$$\Sigma(X) = \frac{\alpha\beta}{2} M^{-1} \tag{89}$$

The following integral

$$\int_{\mathcal{Y}} \exp\left(-\frac{1}{\alpha\beta} (x - K\mu - Ly)^T M (x - K\mu - Ly)\right) dy = \frac{2}{\alpha\beta} (2\pi)^{-\frac{n}{2}} |M|^{-1} \tag{90}$$

is constant irrespective of y . Integrating Eq. 86 over x gives the worst-case probability density function $q^*(y)$:

$$\begin{aligned}
q^*(y) &= \int_{\mathcal{V}} dx \frac{p(x) \exp\left(\frac{V(y)}{\alpha} - \frac{c(x,y)}{\alpha\beta}\right)}{\int_{\Omega} \exp\left(\frac{V(y)}{\alpha} - \frac{c(x,y)}{\alpha\beta}\right) dy} \\
&\propto \int_{\mathcal{V}} \exp\left[\frac{1}{\alpha\beta}(K\mu + Ly)^T M(K\mu + Ly) - \frac{1}{\alpha\beta}y^T(B - \beta A)y\right] dx \\
&= \exp\left[\frac{1}{\alpha\beta}\left(\frac{\alpha\beta}{2}M^{-1}\Sigma^{-1}\mu + M^{-1}By\right)^T M\left(\frac{\alpha\beta}{2}M^{-1}\Sigma^{-1}\mu + M^{-1}By\right) - \frac{1}{\alpha\beta}y^T(B - \beta A)y\right] \\
&= \exp\left[\frac{1}{\alpha\beta}\left(\left(\frac{\alpha\beta}{2}\Sigma^{-1}\mu + By\right)^T M^{-1}\left(\frac{\alpha\beta}{2}\Sigma^{-1}\mu + By\right) - y^T(B - \beta A)y\right)\right] \\
&\propto \exp\left[-\frac{1}{2}(y - B(B - \beta A)^{-1}\mu)^T \right. \\
&\quad \left. \left((B - \beta A)^{-1}B\Sigma B(B - \beta A)^{-1} + \frac{\alpha\beta}{2}(B - \beta A)^{-1}\right)^{-1} (y - B(B - \beta A)^{-1}\mu)\right] \tag{91}
\end{aligned}$$

Eq. 91 shows that the worst-case distribution is still multivariate normal. The vector of means and the covariance matrix are given respectively by

$$\mu_W = (B - \beta A)^{-1}B\mu \tag{92}$$

$$\Sigma_W = (B - \beta A)^{-1}B\Sigma B(B - \beta A)^{-1} + \frac{\alpha\beta}{2}(B - A)^{-1} \tag{93}$$

An interesting observation on Eq. 92 is that the worst-case measure can be generated by a measure-preserving linear map. In fact, for any vector v of asset returns, the linear map g gives

$$\begin{aligned}
g(v) &= (B - \beta A)^{-1}Bv \\
&= (I - \beta B^{-1}A)^{-1}v \tag{94}
\end{aligned}$$

We write down the probability density function for the nominal measure by

$$f(v) \propto \exp\left(-\frac{1}{2}(v - \mu)^T \Sigma^{-1}(v - \mu)\right) \tag{95}$$

The measure given by the measure-preserving map g has a probability density function that is proportional to $f(g^{-1}(v))$,

$$\begin{aligned}
&f(g^{-1}(v)) \\
&\propto \exp\left(-\frac{1}{2}((I - \beta B^{-1}A)v - \mu)^T \Sigma^{-1}((I - \beta B^{-1}A)v - \mu)\right) \\
&= \exp\left(-\frac{1}{2}(v - (I - \beta B^{-1}A)^{-1}\mu)^T (I - \beta B^{-1}A)\Sigma^{-1}(I - \beta B^{-1}A)((I - \beta B^{-1}A)v - \mu)\right) \\
&= \exp\left(-\frac{1}{2}(v - \tilde{\mu})^T \tilde{\Sigma}^{-1}(v - \tilde{\mu})\right) \tag{96}
\end{aligned}$$

where

$$\tilde{\mu} := (I - \beta B^{-1}A)^{-1} \mu \quad (97)$$

$$\tilde{\Sigma} := (I - \beta B^{-1}A)^{-1} \Sigma (I - \beta B^{-1}A)^{-1} \quad (98)$$

that are precisely the mean and covariance matrix given in Eq. 92 (with $\alpha = 0$). As a result, we generate the worst-case measure by applying the measure-preserving map g .

6.5 E. The support of a multivariate normal distribution

In this section, we discuss the support of the nominal measure P assuming the asset returns follow a multivariate normal distribution. In addition, we want see how it is altered by different approaches to model risk measurement. Apparently, approaches based on f -divergence cannot alter the support as they only account for measures that are equivalent to the nominal one. But this conclusion does not tell us explicitly what the support is. In the following work we aim to find the linear subspace that supports the measure.

Formally speaking, returns of the n assets form a n -dimensional vector that lives in a n -dimensional topological vector space \mathcal{V} . If the asset returns follow a multivariate normal distribution with a non-singular covariance matrix, then the support is the entire space \mathcal{V} . However, if the covariance matrix is singular, the support can only be part of \mathcal{V} . We will find this support and show that it is a m -dimensional linear subspace, where m is the rank of the covariance matrix.

The nominal model of asset returns defines a probability space $(\mathcal{V}, \mathcal{F}, P)$, where \mathcal{F} is the Borel σ -algebra on \mathcal{V} . Since \mathcal{V} is a vector space, we may consider its dual space \mathcal{V}^* , i.e. the space of linear maps $a : \mathcal{V} \rightarrow \mathbb{R}$. Any element of the dual space is regarded as a vector of portfolio weights. To see this, suppose the asset returns are $v = (v_1, v_2, \dots, v_n) \in \mathcal{V}$, and the portfolio weights are $a = (a_1, a_2, \dots, a_n) \in \mathcal{V}^*$. The pairing of a and v results in a real number, which is exactly the portfolio return:

$$a(v) = \sum_{j=1}^n a_j v_j \quad (99)$$

If we treat the asset returns v_i as random variables, we may calculate of portfolio variance on a given vector of weights $a \in \mathcal{V}$ by $\text{Var}(a(v)) = a^T \Sigma a$, where Σ is the covariance matrix of the asset returns. For convenience, we use the same symbol v for both the vector of random variables (random vector) and its realization (i.e. a specific element in \mathcal{V}). Now take the positive semi-definite matrix Σ as a linear map $\Sigma : \mathcal{V}^* \rightarrow \mathcal{V}$:

$$\Sigma(a) = \Sigma a \in \mathcal{V}, \quad \forall a \in \mathcal{V}^* \quad (100)$$

The portfolio variance is formed by applying the linear map $a : \mathcal{V} \rightarrow \mathbb{R}$ to $\Sigma(a) \in \mathcal{V}$: $\text{Var}(a(v)) = a(\Sigma(a))$. If the square matrix Σ is singular, then its kernel $\ker \Sigma$ is not trivial (i.e. contains elements other than the zero vector). \mathcal{V}^* can therefore be decomposed into two subspaces:

$$\mathcal{V}^* = \ker \Sigma \oplus \ker \Sigma^\perp \quad (101)$$

Suppose $\ker \Sigma^\perp$ has dimension n . $\ker \Sigma$ has dimension $m - n$ for the dimensions of subspaces sum up to the dimension of \mathcal{V}^* . We may switch to a new orthonormal basis $\{e_1^*, e_2^*, \dots, e_m^*, k_1^*, k_2^*, \dots, k_{m-n}^*\}$ in consistency with the decomposition Eq. 101, in the sense that $e_1^*, e_2^*, \dots, e_m^*$ span $\ker \Sigma^\perp$ and $k_1^*, k_2^*, \dots, k_{m-n}^*$ span $\ker \Sigma$. Now get back the

original space of asset returns \mathcal{V} , we may select a new basis $\{e_1, e_2, \dots, e_m, k_1, k_2, \dots, k_{m-n}\}$ that is dual to $\{e_1^*, e_2^*, \dots, e_m^*, k_1^*, k_2^*, \dots, k_{m-n}^*\}$, i.e.

$$e_i^*(e_j) = \delta_{i-j} \quad (102)$$

$$k_i^*(k_j) = \delta_{i-j} \quad (103)$$

$$e_i^*(k_j) = 0 \quad (104)$$

$$k_i^*(e_j) = 0 \quad (105)$$

Any $v \in \mathcal{V}$ can be expressed by

$$v = \sum_{i=1}^m u_i e_i + \sum_{i=1}^{m-n} w_i k_i \quad (106)$$

Suppose \mathcal{U} denotes the linear subspace spanned by e_1, e_2, \dots, e_m . \mathcal{U} is in fact the dual space of $\ker \Sigma^\perp$. We will show that the support of the nominal measure P is indeed the linear subspace \mathcal{U} shifted by the vector of average asset returns μ :

Theorem Given a finite-dimensional topological vector space \mathcal{V} and its Borel α -algebra \mathcal{F} , the support of a measure P on $(\mathcal{V}, \mathcal{F})$ is $\{v \in \mathcal{V} : v - \mu \in \mathcal{U}\}$ if P provides a multivariate distribution $\mathcal{N}(\mu, \Sigma)$.

Proof For every $v \in \ker \Sigma$, consider the variance of $a(v)$ (v is a random vector here):

$$\text{Var}(a(v)) = a^T \Sigma a = 0 \quad (107)$$

The zero variance implies that a carries the measure P on \mathcal{V} to a Dirac measure P_a on \mathbb{R}

$$P_a(A) = P(a^{-1}(A)), \forall A \in \{A \subseteq \mathbb{R} : a^{-1}(A) \in \mathcal{F}\} \quad (108)$$

Suppose $\text{supp}(P_a) = \{s_a\}$ where $s_a \in \mathbb{R}$. We can show that $\text{supp}(P)$ should only include elements in \mathcal{V} that is projected to s_a . More formally, with the projection map $\mathcal{P} : \mathcal{V} \rightarrow \ker \Sigma$, we have

$$\{v \in \mathcal{V} : \exists a \in \ker \Sigma, a(v) \neq s_a\} \cap \text{supp}(P) = \emptyset \quad (109)$$

In fact, for a given $v \in \mathcal{V}$, suppose there exists $a \in \ker \Sigma$ such that $a(v) \neq s_a$. $a(v)$ is not in the $\text{supp}(P)$, suggesting the existence of an open neighborhood $N_{a(v)} \subseteq \mathbb{R}$ such that $P_a(N_{a(v)}) = 0$. Since the linear map a is continuous, $a^{-1}(N_{a(v)})$ is an open neighborhood of v and

$$P(a^{-1}(N_{a(v)})) = P_a(N_{a(v)}) = 0 \quad (110)$$

As a result, $v \notin \text{supp}(P)$ which proves Eq. 109.

Now we consider the set $S := \{v \in \mathcal{V} : a(v) = s_a, \forall a \in \ker \Sigma\}$. For a given $v_s \in S$, every $v \in S$ satisfies

$$a(v - v_s) = a(v) - a(v_s) = 0, \forall a \in \ker \Sigma \quad (111)$$

suggesting that $v - v_s \in \mathcal{U}$. Therefore $S = \{v \in \mathcal{V} : v - v_s \in \mathcal{U}\}$. Regard \mathcal{U} as a topological linear subspace of \mathcal{V} equipped with the relative topology. $\tilde{\mathcal{F}}$ is the Borel

σ -algebra on \mathcal{U} . We may define a new probability space $(\mathcal{U}, \tilde{\mathcal{F}}, \tilde{P})$ by

$$\tilde{P}(A \cap \mathcal{U}) = P(A), \quad \forall A \in \mathcal{F} \quad (112)$$

One can verify that this probability space is well defined. Now we would like to show that $\text{supp}(\tilde{P}) = \mathcal{U}$. In fact, assuming this is true, then for arbitrary $v \in S$ every open neighborhood $N(v)$ has positive measure:

$$P(N(v)) = \tilde{P}(N(v) \cap \mathcal{U}) > 0 \quad (113)$$

This immediately leads to the result $\text{supp}(P) = S$. In particular, from the property of the multivariate normal distribution, $\text{supp}(P)$ includes the vector μ of average asset returns. This means that $\mu \in S$, and thus the support of P can be written as $\text{supp}(P) = S = \{v \in \mathcal{V} : v - \mu \in \mathcal{U}\}$.

Now we only need to show that $\text{supp}(\tilde{P}) = \mathcal{U}$. Consider the projection map $\mathcal{P} : \mathcal{V} \rightarrow \mathcal{U}$ that sends $v = (u_1, u_2, \dots, u_m, w_1, w_2, \dots, w_{m-n})$ to $u = (u_1, u_2, \dots, u_m)$. The projection results in the marginal distribution w.r.t u_1, u_2, \dots, u_m . This marginal distribution characterises a measure P' on the subspace \mathcal{U} :

$$P'(A) = P(\mathcal{P}^{-1}(A)), \quad \forall A \in \{A \subseteq \mathcal{U} : \mathcal{P}^{-1}(A) \in \mathcal{F}\} \quad (114)$$

For any $A \in \mathcal{F}$,

$$\begin{aligned} P'(A \cap \mathcal{U}) &= P(\mathcal{P}^{-1}(A \cap \mathcal{U})) \\ &= \tilde{P}(\mathcal{P}^{-1}(A \cap \mathcal{U}) \cap \mathcal{U}) \\ &= \tilde{P}(A \cap \mathcal{U}) \end{aligned} \quad (115)$$

Therefore, the two measures \tilde{P} and P' coincide, and we only need to prove $\text{supp}(P') = \mathcal{U}$. The marginal distribution from projection \mathcal{P} is apparently multivariate normal (every linear combination of the elements in u is also a linear combination of the elements in $v \in \mathcal{P}^{-1}(u)$ thus normally distributed).

The covariance matrix $\tilde{\Sigma}$ of the truncated random vector u is invertible. In fact, because $\Sigma(a)$ is not a zero vector for every non-zero $a \in \ker \Sigma^\perp$, the linear map between two m -dimensional vector spaces $\Sigma|_{\ker \Sigma^\perp} : \ker \Sigma^\perp \rightarrow \Sigma(\ker \Sigma^\perp)$ is invertible. Represented by a $m \times m$ matrix, $\Sigma|_{\ker \Sigma^\perp}$ has only non-zero eigenvalues. Since it is also positive semi-definite (for $\text{Var}(a(v)) = a(\Sigma(a)) \geq 0, \forall a \in \ker \Sigma^\perp \subseteq \mathcal{V}^*$), it must be positive definite. We conclude that for every non-zero $a \in \ker \Sigma^\perp$, $\text{Var}(a(v)) = a(\Sigma|_{\ker \Sigma^\perp}(a)) > 0$. If we expand $a(v)$ component-wise according to Eq. 106,

$$\begin{aligned} a(v) &= \sum_{i=1}^m u_i a(e_i) + \sum_{i=1}^{m-n} w_i a(k_i) \\ &= a \left(\sum_{i=1}^m u_i e_i \right) = a(u) \end{aligned} \quad (116)$$

Therefore $a^T \tilde{\Sigma} a = \text{Var}(a(u)) = \text{Var}(a(v)) > 0$ for every nonzero $a \in \ker \Sigma^\perp$. As a result, $\tilde{\Sigma}$ is positive-definite and thus invertible. Under the measure P' , the random vector u follows a multivariate normal distribution with a non-singular covariance matrix. It is supported by the entire subspace \mathcal{U} , i.e. $\text{supp}(P') = \mathcal{U}$. \square

For a multivariate distribution $\mathcal{N}(\mu, \Sigma)$, the support $\text{supp}(P) = \{v \in \mathcal{V} : v - \mu \in \mathcal{U}\}$ only depends on the vector μ and the kernel of Σ . It is clear that under the Kullback-

Leibler divergence the worst-case measure shares the same support. In fact, the worst-case distribution is $\mathcal{N}(\mu_{KL}, \Sigma_{KL})$ where μ_{KL} and Σ_{KL} are given in Eq. 46. Assuming θ is sufficiently small so that $I - 2\theta\Sigma A$ is invertible, $\Sigma_{KL}a = 0$ if and only if $\Sigma_{KL} = 0$ for every $a \in \mathcal{V}^*$. Therefore, Σ_{KL} and Σ share the same kernel and therefore the same subspace $\mathcal{U} \subseteq \mathcal{V}$. In addition, $\mu_{KL} - \mu \in \mathcal{U}$ because for every $a \in \ker\Sigma$ we have

$$\begin{aligned} a(\mu_{KL} - \mu) &= a(2\theta\Sigma A(I - 2\theta\Sigma A)^{-1}\mu) \\ &= 2\theta(\Sigma a)^T A(I - 2\theta\Sigma A)^{-1}\mu \\ &= 0 \end{aligned} \tag{117}$$

As a result, the support of the worst-case measure is $\{v \in \mathcal{V} : v - \mu_{KL} \in \mathcal{U}\} = \{v \in \mathcal{V} : v - \mu \in \mathcal{U}\}$, same as the support of the nominal measure.

On the other hand, the worst-case measured resulted from the Wasserstein approach can have different support. According to Eq. 44, the worst-case covariance matrix Σ_W has a different kernel in general. In addition, $\mu_W - \mu = \beta A(B - \beta A)^{-1}\mu$ is not dependent on Σ , thus not linked to the subspace \mathcal{U} . Setting $\alpha = 0$ in Eq. 44 provides a particularly interesting case, where the worst-case measure is given by a measure-preserving linear map $g : \mathcal{V} \rightarrow \mathcal{V}$ given by Eq. 94. As a result, the support of the worst-case measure can be obtained using the same map, i.e.

$$\begin{aligned} &\{v \in \mathcal{V} : g^{-1}(v) - \mu \in \mathcal{U}\} \\ &= \{v \in \mathcal{V} : (I - \beta B^{-1}A)v - \mu \in \mathcal{U}\} \\ &= \{v \in \mathcal{V} : v - (I - \beta B^{-1}A)^{-1}\mu \in \{(I - \beta B^{-1}A)^{-1}u : u \in \mathcal{U}\}\} \\ &= \{v \in \mathcal{V} : v - \mu_W \in \mathcal{U}_W\} \end{aligned} \tag{118}$$

$\mathcal{U}_W := \{(I - \beta B^{-1}A)^{-1}u : u \in \mathcal{U}\} \subseteq \mathcal{V}$ is the linear subspace (perpendicular to $\ker\Sigma_W$) that corresponds to the worst-case scenario under the Wasserstein approach.

6.6 F. Verification of the Wasserstein approach

Sec. 6.5 shows that under the Wasserstein approach the worst-case measure does alter the support. Now the question is whether the approach searches over all alternative measures. Unlike f -divergence that is only capable of measuring distance between equivalent measures, the Wasserstein metric provides a finite distance between non-equivalent measures as well. Therefore the Wasserstein approach should be able to find out the worst-case measure from all equivalent and non-equivalent measures. In this section, we will verify it for the example of portfolio variance. In particular, we will find out a worst-case linear map $g^* : \mathcal{V} \rightarrow \mathcal{V}$ by searching over the entire space of linear maps. We will verify that Eq. 46 (with $\alpha = 0$) can be given by the worst-case linear map.

Theorem Given a probability space $(\mathcal{V}, \mathcal{F}, P)$ where \mathcal{V} is a finite-dimensional vector space and P provides a multivariate distribution $\mathcal{N}(\mu, \Sigma)$, there exists a worst-case linear map $g^* : \mathcal{V} \rightarrow \mathcal{V}$ in the sense of Eq. 15, i.e.

$$g^*(x) = \arg \max_{y \in \mathcal{V}} \left[y^T A y - \frac{(x - y)^T B (x - y)}{\beta} \right] \tag{119}$$

for every non-zero $x \in \mathcal{V}$, as long as $B - \beta A$ is positive definite.

Proof Given a non-zero $x \in \mathcal{V}$, every non-zero $y \in \mathcal{V}$ can be expressed by $y = g(x)$ where g is some linear map (not unique) $g : \mathcal{V} \rightarrow \mathcal{V}$. The problem Eq. 119 is therefore equivalent

to

$$g^*(x) = \arg \max_{g \in \mathfrak{L}(\mathcal{V}, \mathcal{V})} \left[g(x)^T A g(x) - \frac{(x - g(x))^T B (x - g(x))}{\beta} \right] (x) \quad (120)$$

where $\mathfrak{L}(\mathcal{V}, \mathcal{V})$ is the space of all linear maps from \mathcal{V} to \mathcal{V} . Choosing an orthonormal basis for \mathcal{V} allows us to represent g by a square matrix, and the linear map $g(x)$ by matrix multiplication gx . The expression inside the square bracket in Eq. 120 is then transformed into

$$\begin{aligned} & (gx)^T A gx - \frac{(x - gx)^T B (x - gx)}{\beta} \\ &= -\frac{1}{\beta} x^T (g^T - B(B - \beta A)^{-1}) (B - \beta A) (g - (B - \beta A)^{-1} B) x \\ & \quad - \frac{1}{\beta} x^T (B - B(B - \beta A)^{-1} B) x \end{aligned} \quad (121)$$

Since $B - \beta A$ is positive-definite, the first term in Eq. 121 is either zero or negative. It reaches zero (and hence Eq. 121 reaches its maximum value) if and only if

$$(g - (B - \beta A)^{-1} B) x = 0 \quad (122)$$

or equivalently

$$g(x) = (B - \beta A)^{-1} B x \quad (123)$$

This allows to rewrite the right-hand side of Eq. 120:

$$g^*(x) = (B - \beta A)^{-1} B x \quad (124)$$

The linear map g^* given by the square matrix $(B - \beta A)^{-1} B$ satisfies Eq. 124 for every non-zero $x \in \mathcal{V}$. \square

It is noted that in the problem of portfolio variance risk, both square matrices A and B are symmetric and positive definite. Therefore if the positive multiplier β is sufficiently small, $B - \beta A$ is also positive definite satisfying the condition assumed by the theorem above. Now the worst-case linear map g^* transforms the vector of asset returns from μ to $(B - \beta A)^{-1} B \mu$, and the covariance matrix from Σ to $(B - \beta A)^{-1} B \Sigma B (B - \beta A)^{-1}$, same as the expressions given in Eq. 46 (with $\alpha = 0$). This verifies that the Wasserstein approach indeed searches over the entire space $\mathfrak{L}(\mathcal{V}, \mathcal{V})$ of linear maps. It results in a measure that corresponds to the worst-case linear map g^* .

6.7 G. Robust MVO Portfolio (Kullback-Leibler divergence)

According to Eq. 54, we consider the problem

$$\max_{Q \in \mathcal{M}} E^Q \left((X - \mu - k)^T a a^T (X - \mu - k) \right) \quad (125)$$

Since $X - \mu - k \sim \mathcal{N}(-k, \Sigma)$, under the Kullback-Leibler divergence the covariance matrix and the mean of the worst-case measure are given according to Eq. 46 (remember

$a^T k = \lambda/2$):

$$\begin{aligned}\Sigma_{KL} &= (I - 2\theta\Sigma A)^{-1}\Sigma & (126) \\ \mu_{KL} &= (\mu + k) - (I - 2\theta\Sigma A)^{-1}k \\ &= \mu - \lambda\theta(I - 2\theta\Sigma A)^{-1}\Sigma a\end{aligned}$$

Using the Wassertein approach, however, the worst-case measure has different covariance matrix and mean (Eq. 44):

$$\begin{aligned}\Sigma_W &= (I - \beta B^{-1}A)^{-1}\Sigma(I - \beta AB^{-1})^{-1} & (127) \\ \mu_W &= (\mu + k) - (I - \beta B^{-1}A)^{-1}k \\ &= \mu - \frac{\lambda}{2}\beta(I - \beta B^{-1}A)^{-1}B^{-1}a\end{aligned}$$

We may then formulate the optimal asset allocation a^* under the worst-case measure. According to Eq. 52, the problem is formulated in the following form under the Kullback-Leibler divergence.

$$\begin{aligned}\min_a & a^T \Sigma_{KL} a - \lambda a^T \mu_{KL} \\ &= a^T (I - 2\theta\Sigma A)^{-1} \Sigma a - \lambda a^T (\mu - \lambda\theta(I - 2\theta\Sigma A)^{-1} \Sigma a) \\ &= a^T \Sigma a - \lambda a^T \mu + \theta a^T \Sigma a (\lambda^2 + 2a^T \Sigma a) + O(\theta^2)\end{aligned} \quad (128)$$

To find out a closed-form solution, we need to ignore the higher order terms $O(\theta^2)$. Then the stationary condition of the minimisation problem is given by a non-linear equation:

$$2\Sigma a - \lambda\mu + 2\theta(\lambda^2 + 4a^T \Sigma a)\Sigma a = 0 \quad (129)$$

Notice that

$$a^* := \frac{\lambda}{2}\Sigma^{-1}\mu \quad (130)$$

is the MVO portfolio weight under the nominal measure. For the robust MVO portfolio, we may consider its first-order deviation from a^* . To do that, we substitute $a = a^* + \theta b$ into Eq. 129 allowing us to cancel the term $\lambda\mu$.

$$2\theta\Sigma b + \theta\lambda(\lambda^2 + \lambda^2\mu^T \Sigma^{-1}\mu)\mu + O(\theta^2) = 0 \quad (131)$$

By matching the first-order term w.r.t θ , we find the expression for b :

$$b = -\frac{\lambda^3}{2}(1 + \mu^T \Sigma^{-1}\mu)\Sigma^{-1}\mu \quad (132)$$

Therefore the optimal MVO portfolio under the worst-case scenario is

$$\begin{aligned}a_{KL}^* &= \left(\frac{\lambda}{2} - \frac{\theta\lambda^3}{2}(1 + \mu^T \Sigma^{-1}\mu)\right)\Sigma^{-1}\mu \\ &= ca^*\end{aligned} \quad (133)$$

where the coefficient c is defined by

$$c := 1 - \theta\lambda^2(1 + \mu^T \Sigma^{-1}\mu) \quad (134)$$

The robust MVO portfolio, as a vector a_{KL}^* , is parallel to the normal MVO portfolio a^* . As a result, the robust MVO portfolio does not change the relative weights of component assets. In fact, all the weights are reduced by the same proportion ($c < 1$) to account for model risk. This is, however, inappropriately account for the correlation risk. For example, two highly-correlated assets have extremely high weights in the nominal MVO portfolio. Because of the correlation risk, we would expect the robust MVO portfolio to assign them lower weights relative to other assets.

The Sharpe ratio of the robust MVO portfolio obviously equals the Sharpe ratio under the nominal measure, denoted by S ($S = \sqrt{\mu^T \Sigma^{-1} \mu}$). Sometimes we may be interested in the Sharpe ratio under the worst-case measure. This requires us to examine the mean and variance of the robust MVO portfolio given by Eq. 133. Assuming that we are under the worst-case scenario given by Eq. 126, the portfolio mean and variance can be obtained by substituting $a_{KL}^* = c\lambda \Sigma^{-1} \mu / 2$:

$$\mu_{KL}^T a_{KL}^* = \left(\mu - \lambda \theta (I - 2\theta \Sigma a_{KL}^* a_{KL}^{*T})^{-1} \Sigma a_{KL}^* \right)^T a_{KL}^* \quad (135)$$

$$= \frac{\lambda c}{2} \mu^T \Sigma^{-1} \mu - \theta \frac{\lambda^3 c^2}{4} \mu^T \Sigma^{-1} \mu + O(\theta^2) \quad (136)$$

$$a_{KL}^{*T} \Sigma_{KL} a_{KL}^* = a_{KL}^{*T} (I - 2\theta \Sigma a_{KL}^* a_{KL}^{*T})^{-1} \Sigma a_{KL}^* \quad (137)$$

$$= \frac{\lambda^2 c^2}{4} \mu^T \Sigma^{-1} \mu + \theta \frac{\lambda^4 c^4}{8} (\mu^T \Sigma^{-1} \mu)^2 + O(\theta^2) \quad (138)$$

By using the portfolio mean and variance given in Eq. 135, we may calculate the Sharpe ratio of the robust MVO portfolio (under the worst-case scenario):

$$\begin{aligned} S_{KL} &= \frac{\mu_{KL}^T a_{KL}^*}{\sqrt{a_{KL}^{*T} \Sigma_{KL} a_{KL}^*}} \\ &= \frac{2 - \theta \lambda^2 c + O(\theta^2)}{2 + \frac{\theta}{2} \lambda^2 c^2 \mu^T \Sigma^{-1} \mu + O(\theta^2)} \sqrt{\mu^T \Sigma^{-1} \mu} \\ &= \left(1 - \frac{\theta}{4} \lambda^2 c (cS^2 + 2) + O(\theta^2) \right) S \end{aligned} \quad (139)$$

We can see that the robust Sharpe ratio (defined as the Sharpe ratio of the robust MVO portfolio under the worst-case model) is a function of the nominal Sharpe ratio S . The MVO portfolio corresponds to $c = 1$, suffering from more reduction in Sharpe ratio than the robust MVO portfolio ($c < 1$) under the worst-case measure. This simple relation, however, no longer holds for the Wasserstein approach.

6.8 H. Robust MVO Portfolio (Wasserstein approach)

In this section, we will switch to the Wasserstein approach to model risk measurement. We will derive the robust MVO portfolio with the Wasserstein approach. Using Eq. 127, we may formulate the robust portfolio optimisation problem in the following form:

$$\begin{aligned} \min_a \quad & a^T \Sigma_W a - \lambda a^T \mu_W \\ &= a^T (I - \beta B^{-1} A)^{-1} \Sigma (I - \beta A B^{-1})^{-1} a - \lambda a^T \left(\mu - \frac{\lambda}{2} \beta (I - \beta B^{-1} A)^{-1} B^{-1} a \right) \\ &= a^T \Sigma a - \lambda a^T \mu + \beta \left(2a^T B^{-1} a a^T \Sigma a + \frac{\lambda^2}{2} a^T B^{-1} a \right) + O(\beta^2) \end{aligned} \quad (140)$$

Ignoring the higher order terms, the minimisation problem is solved using

$$2\Sigma a - \lambda\mu + \beta (4a^T B^{-1} a \Sigma a + (4a^T \Sigma a + \lambda^2) B^{-1} a) = 0 \quad (141)$$

Substituting $a = a^* + \beta b$ into Eq. 141, we find the expression for the perturbation b by matching the first-order terms of β :

$$b = -\frac{\lambda^3}{4} (\mu^T \Sigma^{-1} B^{-1} \Sigma^{-1} \mu + (1 + \mu^T \Sigma^{-1} \mu) \Sigma^{-1} B^{-1}) \Sigma^{-1} \mu \quad (142)$$

Therefore, the Wasserstein approach results in a robust MVO portfolio with weights

$$\begin{aligned} a_W^* &= \left(\frac{\lambda}{2} - \frac{\beta\lambda^3}{4} \mu^T \Sigma^{-1} B^{-1} \Sigma^{-1} \mu - \frac{\beta\lambda^3}{4} (1 + \mu^T \Sigma^{-1} \mu) \Sigma^{-1} B^{-1} \right) \Sigma^{-1} \mu \\ &= c a^* - D a^* \end{aligned} \quad (143)$$

where c is a coefficient while D is a square matrix defined by

$$\begin{aligned} c &:= 1 - \frac{\beta\lambda^2}{2} \mu^T \Sigma^{-1} B^{-1} \Sigma^{-1} \mu \\ D &:= \frac{\beta\lambda^2}{2} (1 + \mu^T \Sigma^{-1} \mu) \Sigma^{-1} B^{-1} := d \Sigma^{-1} B^{-1} \end{aligned} \quad (144)$$

c serves just as the coefficient under the Kullback-Leibler divergence, reducing the portfolio weights by the same fraction. D is a matrix that serves to linearly transform the normal MVO portfolio weights.

Eq. 143 correctly accounts for the correlation risk. When two assets are highly correlated, Σ is close to be singular. This results in extremely large weights under the normal MVO portfolio. Eq. 143, on the other hand, not only scales the weights down simultaneously by the coefficient c , but also reduces the relative weights of the highly-correlated assets by the linear map D . To see how the linear map D changes the relative weights, we may re-arrange Eq. 143 to the following form:

$$a_W^* = \frac{\lambda}{2} (\Sigma(cI - D)^{-1})^{-1} \mu \quad (145)$$

Therefore, the robust MVO portfolio has the same weights as a normal MVO portfolio with an effective covariance matrix

$$\Sigma^* = \Sigma(cI - d\Sigma^{-1}B^{-1})^{-1} \quad (146)$$

One can show by induction that $\Sigma v = xv$ (x and v are respectively the eigenvalue and the eigenvector) leads to $\Sigma^n v = xv$ for every integer n . This is to say, x is an eigenvalue of Σ only if x^n is an eigenvalue of Σ^n corresponding to the same eigenvector. As a result, for every eigenvalue $x > 0$ of the positive-definite covariance matrix, there exists a corresponding eigenvalue of the effective covariance matrix (here we only consider the special

case where B is the identity matrix I):

$$\begin{aligned}
\Sigma^* v &= \Sigma(cI - d\Sigma^{-1})^{-1} v \\
&= \frac{1}{c} \sum_{i=0}^{\infty} \left(\frac{d}{c}\right)^i \Sigma^{1-i} v \\
&= \frac{1}{c} \sum_{i=0}^{\infty} \left(\frac{d}{c}\right)^i x^{1-i} v \\
&= \frac{x}{c - d/x} v
\end{aligned} \tag{147}$$

The corresponding eigenvalue

$$\begin{aligned}
x^* &:= \frac{x}{c - d/x} \\
&= x + \frac{\beta\lambda^2}{2} (1 + \mu^T \Sigma^{-1} \mu + \mu^T \Sigma^{-1} B^{-1} \Sigma^{-1} \mu) + O(\beta^2)
\end{aligned} \tag{148}$$

Any eigenvalue x close to zero is adjusted according to Eq. 148, resulting in a corresponding eigenvalue x^* that is at least as large as $\beta\lambda^2/2$. This results in an effective matrix Σ^* that is less "singular" than Σ , and therefore a robust MVO portfolio that accounts for the correlation risk.




Fig 11: Eigenvalue x^* of the effective covariance matrix Σ^* increases by a greater amount when the original eigenvalue x gets closer to zero.

References

- [1] Syed Mumtaz Ali and Samuel D Silvey. A general class of coefficients of divergence of one distribution from another. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 131–142, 1966.
- [2] I Csisz et al. Information-type measures of difference of probability distributions and indirect observations. *Studia Sci. Math. Hungar.*, 2:299–318, 1967.
- [3] Amir Ahmadi-Javid. Entropic value-at-risk: A new coherent risk measure. *Journal of Optimization Theory and Applications*, 155(3):1105–1123, 2012.

- [4] Paul Glasserman and Xingbo Xu. Robust risk measurement and model risk. *Quantitative Finance*, 14(1):29–58, 2014.
- [5] C Villani. *Optimal transport – Old and new*, volume 338, page 18. 01 2008.
- [6] C Villani. *Optimal transport – Old and new*, volume 338, page 22. 01 2008.
- [7] Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- [8] Gurdip Bakshi and Nikunj Kapadia. Delta-hedged gains and the negative market volatility risk premium. *The Review of Financial Studies*, 16(2):527–566, 2003.
- [9] Buen Sin Low and Shaojun Zhang. The volatility risk premium embedded in currency options. *Journal of Financial and Quantitative Analysis*, 40(4):803–832, 2005.
- [10] Peter Carr and Liuren Wu. Analyzing volatility risk and risk premium in option contracts: A new theory. *Journal of Financial Economics*, 120(1):1–20, 2016.
- [11] Tim Bollerslev, George Tauchen, and Hao Zhou. Expected stock returns and variance risk premia. *The Review of Financial Studies*, 22(11):4463–4492, 2009.
- [12] Yacine Ait-Sahalia, Mustafa Karaman, and Lorian Mancini. The term structure of variance swaps and risk premia. 2015.