13th International Conference on Applications of Statistics and Probability in Civil Engineering, ICASP13
Seoul, South Korea, May 26-30, 2019

Tuning deep flooding risk with adaptive strategy: An application for NYC

Kairui Feng
Graduate Student, Dept. of Civil and Environmental Engineering, Princeton University, Princeton,U.S.A

Siyuan Xian
Ph.D, Dept. of Civil and Environmental Engineering, Princeton University, Princeton,U.S.A

Ning Lin
Associate Professor, Dept. of Civil and Environmental Engineering, Princeton University, Princeton,U.S.A

ABSTRACT: The impact of the storms may worsen in the coming decades due to the rapid development of the coastal zone in conjunction with sea-level rise and possibly increased storm activity due to climate change. Greater progress on coastal flood risk management is urgently needed. Previous studies proposed designs of dynamic seawalls (i.e., seawalls that can be heightened overtime to cope with the increasing effect of climate change), based on long-term climate model projections. However, significant uncertainties exist in long-term climate projections. Noticing that the climate condition can be observed over time, we develop a reinforcement-learning-based strategy of adaptive seawall design (i.e., the design is planned to be regularly updated based on observations), to cope with the deep uncertainty in climate change effects. We apply this method to New York City and estimate its optimal adaptive seawall design, based on climate projections of sea-level rise and storm surge flooding, building level exposure data, and estimated construction cost of the seawall. We show that the total lifetime cost (including the investment of the seawall and potential damage of the protected area) is significantly reduced (by 20% to 40%) when the dynamic, reinforced learning strategy is applied, compared to traditional design methods.

1. INTRODUCTION

Coastal flood risk will likely increase in the future due to urban development, sea-level rise, subsidence and potential change of storm surge climatology. Many coastal cities are proceeding to plan for increasing flood risks, as are government agencies and private firms with vulnerable facilities for lower Manhattan. One of the most typical examples is the Big U plan– a $836-million 15-feet seawall motivated by Hurricane Sandy(2012) above the mean sea level – to protect the $500 billion business sector that influences the world's economy (NYCGovernment (2017)). Recent Hurricane Harvey (2017),Irma (2017) and Maria (2017) also arouse the risk awareness of the local governments (Shultz and Galea (2017); Xian et al. (2018)).

Most local governments are considering a static coastal protection plan, such as Big U (New York) or Galveston Seawall (Texas) because of the tradition of static infrastructure design. Aerts et al. (2014) also discussed about vairous static plans for NYC.The static method is rooted in the long history of flood-zoning and mapping. These maps developed for risk calculations — such as FEMA's

HAZUS system for estimating the potential losses from disasters — are also used to guide the protection investments. Protection from events with a particular return period, such as the 100/500-year flood (usually used for flood maps) under the current climate, may be augmented by judgmental safety factors to account for the effect of climate change.

However, the nature of climate change is time-varying and deeply uncertain over time. The protection decisions made only based on current climate information cannot remain optimal under climate change, which brings increasing risks of rising sea level and destructive hurricane potential (Emanuel (2013); Kopp et al. (2014)). Thus, the adaptive seawall design - allowing seawall height to vary both temporally and spatially - becomes economically attractive. The flexible design can be at least as good as static-optimal design because any static plan is a feasible solution for the adaptive optimization problem (Adam and Smith (2008)). Many efforts have been made on dynamic seawall design based on life-cycle cost (Lickley et al. (2014); van der Pol et al. (2017)).These two studies can represent two views of adaptive seawall design: Dynamic Programming and Bayesian Learning, respectively Lickley et al. (2014) brought forward a viewpoint that given the future risk estimation, the adaptive design can improve the economic performance of the seawall project. However, Lickley et al. (2014) didn't directly address a more significance advantage of the flexible design that we can learn the climate evolution and improve our decisions based on the possible future exogenous information. van der Pol et al. (2017) shows a framework of a multistage method to analysis the Bayesian optimization problem coupled with learning. Every time step, they enumerate future scenarios for several steps and choose the best action for this time step. This method successfully coupled learning into the analysis. However, they did not apply the method to a real example. Moreover, their method can deal with only several time steps because of the exponentially increasing complexity of the forwarding optimization method. Finally, the Bayesian learning is not the extreme of the stochastic optimization for seawall design problem.

The value of information may be represented in three parts:1) projected dynamic climate risk informing a dynamic design,2) continuous observation/policy adjusting the design, and 3) a framework of future learning and updating reducing uncertainties for current decision making. The dynamic programming approach of Lickley et al. (2014) reflects the first type of value of climate information. The Bayesian learning approach of van der Pol et al. (2017) also addressed the second value. But they did not address the third type of value.

In this paper, we develop a reinforcement-learning-based strategy of adaptive seawall design (i.e., the design is planned to be regularly updated based on observations), to cope with the deep uncertainty in climate change effects. Mathematically speaking, dynamic programming regards the future climate change independent of previous observations; Bayesian learning regards the future climate change dependent on what has happened; Reinforcement learning regards the future climate change conditioning on all the information before that specific time point. This method strictly improves both dynamic programming method and Bayesian learning method. We apply this method to New York City and estimate its optimal adaptive seawall design, based on climate projections of sea-level rise and storm surge flooding, building level exposure data, and estimated construction cost of the seawall. We show that the total lifetime cost (including the investment of the seawall and potential damage of the protected area) is significantly reduced when the reinforced learning strategy is applied, compared to traditional design methods. Meanwhile, this method also shows better performance in controlling the risk. Thus, there is a great potential of applying reinforcement learning to tune the large global climate risk by human action. Also, since the risk is well controlled under reinforcement learning framework, people can take more aggressive climate adaptation. These results may stimulate the local governments to design policy against climate change because the risk of a governmental project is more clearly estimated under reinforcement learning and the method can help the local government to save resources.

Although this article is going to discuss the advantages and disadvantages for different optimization methods, based on probability measure and risk-neutral setting. The fuzzy math based robust decision making (e.g. minimax regret, information gap) and the risk-attitude behavior are not discussed here (Haasnoot et al. (2013); Ranger et al. (2013)).These approaches may be compared with the optimization discussed here in a future study.

2. METHODOLOGY

The objective function (expected life-cycle cost for T years with discounting rate r) for seawall height management can be separated into two parts:expected damage ($D(\vec{A}, s_t, p_t)$) for the protected area under given seawall height time series ($\vec{A}$) and climatology change of sea level and storm (sea-level change over time, s_t and annual surge exceedance probability distribution, $p_t(X > x)$); The construction cost ($C(\vec{A})$) is also a function of $\vec{A}$. Considering seawall cannot be upgraded at every arbitrary time point (which is a continuous time optimization problem) in reality, here we assume we upgrade the seawall at the end of every δ years and solve this problem as a discrete time optimization. When δ approaches 0, the problem degenerates to a continuous time problem and thus this discretion treatment does not affect the generality of this problem. Here we assume T is divisible by δ and $\vec{A}$ is a $1 \times k$ vector ($k = T/\delta$, regarding the initial seawall height as A_0 and note $A_{-1} = 0$) . Here we assume the expected damage for the considered area for a specific time (i) is related to current seawall height $A_{[i/\delta]+1}$. The cost of the construction is a function of both $A_{[i/\delta]+1}$ and $A_{[i/\delta]}$. Under these settings, the objective function can mathematically be written as:

$$\begin{aligned} L(\vec{A}_t) &= \int_0^T [D(\vec{A}, s_t, p_t) + C(\vec{A})]e^{-rt}dt \\ &= \sum_{t=1}^{T} [D_t(\vec{A}, s_t, p_t) + C(\vec{A})]e^{-rt} \\ &\approx \sum_{i=0}^{k}\sum_{m=1}^{\delta} [D(A_i, s_{i\delta+m}, p_{i\delta+m}) \\ &+ C(A_i, A_{i-1}, m)]e^{-r(i\delta+m)} \end{aligned} \tag{1}$$

The seawall height should never be decreased:

$$A_{i+1} \geq A_i, \forall i \geq -1 \tag{2}$$

And the construction cost is calculated as a linear function of the seawall increment and only happens at the beginning of each period(C_b is the unit price for seawall ($/mile/ft) and $l(h)$ is the length of coastline below seawall height h):

$$C(A_i, A_{i-1}, m) = \begin{cases} \int_{A_{i-1}}^{A_i} C_b l(h)dh & \text{if } m = 1 \\ 0 & \text{else} \end{cases} \tag{3}$$

Under these setting, we build up the framework for dynamic programming, Bayesian learning and reinforcement learning to solve the adaptive seawall height problem.

2.1. *Dynamic Programming*

Assuming the future is uncertain, the dynamic programming method employs the surge risk modeling under future climate (Lin et al. (2012)) and performs the optimization for the objective function. In this model, Lickley et al. (2014) backwardly solved the problem because the seawall built before one specific time point will not affect the future action. However, the possible future action may affect the seawall height now. Thus, by solving the problem backwardly, we can reach the optimal of seawall design, because the decision we calculate out in former step will not conflict with the decision later. This logic holds for all three methods.

Mathematically, dynamic programming is calculating the following problem in order for t = k,k-1,...,0(when calculating A_t, we are assuming A_{t-1}=0. This is reasonable because of the linearity and time-invariance of seawall cost over height.):

$$\begin{aligned} A_t = \text{argmin}_{A_t} \sum_{i=t}^{k}\sum_{m=1}^{\delta} &[D(A_i, s_{i\delta+m}, p_{i\delta+m}) \\ &+ C(A_i, A_{i-1}, m)]e^{-r(i\delta+m)} \end{aligned} \tag{4}$$

For the i th increment of seawall, it covers the city for δ years and for a specific year y in that period, the damage function can be calculated as (here we are using static mapping method for the inundation process to keep the mathematical model simple):

$$D(A_t, s_y, p_y) = \int_{A_t}^{+\infty} e(t) \cdot d(x) \int_{-\infty}^{+\infty} P_t(x - s_y) f(s_y) ds_y dx \quad (5)$$

In which x is the potential surge height; $e(t)$ is the development projection (future exposure divided by the initial exposure of city); $d(x)$ is the damage estimation for the protected urban area under a given surge height x(which integrates the local building distribution and digital elevation model); P_t is the projected annual surge distribution under future climate; $f(s_y)$ is the distribution of sea level at year y. Here a convolution method is applied to estimate the total risk of storm surge(Lin and Shullman (2017)).

This approach successfully changes the sequential decision-making problem on multi-dimension to several 1-dimension problems at one time. For each 1-dimension problem, a 1-dimension search method can be applied to solve the problem.

2.2. *Bayesian Learning*

Bayesian learning approach here partially followed van der Pol et al. (2017). The Bayesian learning algorithm required a future "observation," for the current stage we can assume our probability projection for future is accurate and apply Monte Carlo sampling to simulate possible future circumstances. Whenever we observe a sea level at specific time y, we refresh our understanding of the climate condition (we assume hurricane condition under climate change is not observable because of the limited data size). Moreover, based on the refreshed information, new dynamic programming is down for the period for y to the end to instead of the biased original dynamic programming.

The first algorithmic difference between dynamic programming is the time window we are solving the dynamic programming: if we have already observed the sea level condition for year y, we need to do the dynamic programming for years post y.

The second difference is we need to use conditional probability to estimate the future sea level change under current condition. This will just affect function $D(A_t, s_y, p_y)$. Assuming we have observed the sea level $s_{1:y_0}$ before year y_0, for any year y later than y_0, we have:

$$D(A_t, s_y, p_y) = \int_{A_t}^{+\infty} e(t) \cdot d(x) \int_{-\infty}^{+\infty} P_t(x - s_y) f(s_y | s_{1:y_0}) ds_y dx \quad (6)$$

The full solution follows the procedure in Algorithm 1.

Algorithm 1 Solving Bayesian Learning Optimization for Sea Wall

1: **for** $i_1 = 0 \rightarrow k-1$ **do**
2: $\quad y_0 = i \cdot \delta$, Observe s_{y_0}
3: $\quad$ **for** $i_2 = k \rightarrow i_1$ **do**
4: $\quad\quad \hat{A}_{i_2} = \text{argmin}_{\hat{A}_{i_2}} \sum_{i=i_2}^{k} \sum_{m=1}^{\delta}$
5: $\quad\quad [D(\hat{A}_{i_2}) + C(\hat{A}_{i_2}, \hat{A}_{i_2-1}, m)] e^{-r(i\delta+m)}$
6: $\quad$ **end for**
7: $\quad$ Make decision $A_{i_1} = \hat{A}_{i_1}$
8: **end for**

2.3. *Reinforcement Learning*

Noticing that actually for any time point y we are going to make a decision, we can assume that we have an observation of sea level s_y. However, this effect is not modeled in Bayesian Learning. In reality, if we know that we can obtain new observations in future when we are making decisions, the estimated risk for future becomes much lower and thus can help us better understand the real risk and give out more economic policy - if the real risk is lower than estimated, we tend to be conservative; now we can make risk-neutral decision.

The reinforcement learning process usually can be separated into two parts:1) optimization schema and 2)value function approximation (Powell (2007)). Optimization schema is the methodology people use to search for the optimal solution, such as Monte Carlo Tree search, dynamic programming, etc. Value function approximation is usually used to estimate the outcome (revenue) of every decision we made. Because of the dimension of decision variables (n steps and m choices at each step make the total possible decision size m^n) and possible observations (all the possible scenarios in future), to calculate out the outcome (revenue) of every decision by the exact optimization schema seems not possible, so people need techniques such as neural network or decision space downscaling to approximate the outcome (revenue) of every decision they made at this time point.

Here we apply backward approximate dynamic programming (BADP) to solve this seawall design problem because the strict optimality of the backward dynamic programming approach has been proved in section 2.1 & 2.2. Moreover, also, the approximated decision and observation space compact the variable space considerably then the direct way. Finally, although the feasible space for the stochastic optimization problem in our original version is continuous, by densification of our variable space, our solution can finally approach the exact solution.

The main goal of BADP is to design a look-up table to tell the policymaker what to do when they observed new information. For every possible observation in future, there is an action mapping to that given the table produced by BADP. In this problem, if we know the decision and sea level record for the first $T-\delta$ years ($\tilde{\vec{s}}_{T-\delta}$;here we use $\tilde{x}$ to present a variable is predicted by model to distinguish with the real observation $\hat{x}$), then we can simply solve the final year action A_k:

$$\tilde{A}_k(\tilde{\vec{s}}_{T-\delta}, \tilde{A}_{k-1}) = \arg \min_{\tilde{A}_k \geq \tilde{A}_{k-1}} \sum_{m=1}^{\delta} D(\tilde{A}_k|\tilde{\vec{s}}, \tilde{A}_{k-1}) + C(\tilde{A}_k, \tilde{A}_{k-1}, m) \tag{7}$$

And here we define a look up table $H(\vec{s}_{(i-1)\delta}, A_{i-1}, i)$ to check the action to take based on observed information $\vec{s}_{(i-1)\delta}$ and the i th action. Here for the k th action, we enumerate all the possible $\vec{s}_{(k-1)\delta}$ and made this table by recording the A_k calculated out in eq.7.

Similarly, for other action time steps, we use Algorithm 2 to record them and make the table backwardly.

When we take this algorithm into practice, the probability density function $f(\tilde{\vec{s}}_{i_1+1}|\tilde{\vec{s}}')$ is very difficult to estimate because $\tilde{\vec{s}}'$ is very high-dimensional. So we applied a cluster every $\tilde{\vec{s}}'$ into 3 types: Up, Down and Linear. Then we can use the information of $\tilde{s}_{i-1}$ and the curvature property to project $\tilde{s}_i$.

Algorithm 2 Solving Reinforcement Learning Optimization for Sea Wall

$H(\vec{s}_{(i-1)\delta}, A_{i-1}, i) = [\,]$ // look-up table
$R(\vec{s}_{(i-1)\delta}, A_{i-1}) = [\,]$ // Cumulative damage from step i to k.
for $i_1 = k \to 1$ **do**
    for $\tilde{\vec{s}}'$ *in* $\vec{s}_{(i_1-1)\delta}$
    for $\tilde{A}_{i_1-1}$ *in* All Seawall Height **do**
        $\tilde{A}_{i_1} = \arg\min_{\tilde{A}_{i_1} \geq \tilde{A}_{i_1-1}} \sum_{m=1}^{\delta}$
        $[D(\tilde{A}_{i_1}|\tilde{s}^{i_2}, \tilde{A}_{i_1-1}) + C(\tilde{A}_{i_1}, \tilde{A}_{i_1-1}, m)]+$
        $e^{-r\delta} \int_{-\infty}^{+\infty} f(\tilde{\vec{s}}_{i_1+1}|\tilde{\vec{s}}') R(\tilde{\vec{s}}_{i_1+1}, \tilde{A}_{i_1}) d\tilde{s}_{i_1+1}$
    end for
    Store $H(\vec{s}_{(i-1)\delta}, A_{i-1}, i) = A_i$
    Store $R(\vec{s}_{(i-1)\delta}, A_{i-1})$
    end for
end for

3. EXPERIMENT SETTING

We applied the mentioned methodology to the NYC "Big U" area as shown in Fig. 1.The BIG U calls for a protective system around the low-lying topography of Manhattan beginning at West 57th Street, going down to The Battery, and then back up to East 42nd Street. United States Department of Housing and Urban Development(HUD) has dedicated a total of \$511 million, including Rebuild by Design and National Disaster Resilience Competition funding, toward the implementation of The BIG U, and New York City has committed an additional \$305 million in capital funding to start the first phases of the East Side Coastal Resiliency (ESCR), and Lower Manhattan Coastal Resiliency (LMCR) projects. This project is elegantly designed and considering various functions and areas the city has. However, the height design for the seawall (15 ft above sea level on average) of the "Big U" arouses our interest. This number is taken mainly based on FEMA flood map add 100-year projected sea level with the maximum likelihood. Here we concern whether this simple method can capture the real risk and optimal life-cycle cost of this megaproject. Our design covers 100 years from 2000 to 2100 for the seawall project. The seawall is assumed to be elevated every 10 years.

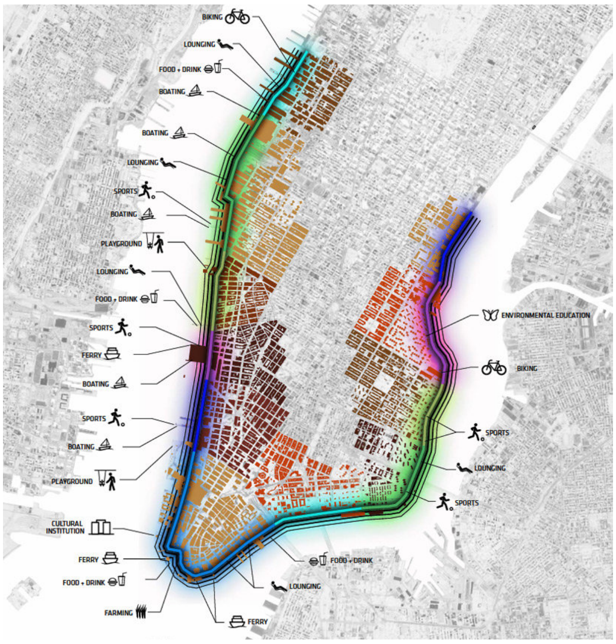

Figure 1: The Big U: a protective system that encircles Manhattan

To estimate the future loss, our sea level scenarios are built on the data set built by Kopp et al. (2014) and Rasmussen et al. (2016). Our annual exceedance probability distribution for storm surge is simulated by Lin et al. (2016). Our building damage setting and seawall construction cost are from Aerts et al. (2014). The discounting rate for seawall construction cost is chosen as 5%.

The total construction cost we estimated for the seawall part of "big U" is about $ 750 million which is quite near the planned budget $ 836 million, as they still have other non-structural systems that also use the money. This allows us to believe the reasonability of our estimation for the construction cost.

4. RESULTS

Different methods suggest different seawall height series as shown in Fig.2. The Big U original design is shown as the green line, which is static at 15 ft. We also based on life-cycle cost and searched one static optimal level for the seawall around Big U area which should be ~17 ft. The bold red line here shows the results of Dynamic Programming. It starts at ∼ 14.5 ft, which is lower than static optimal and afterward it increases quickly to ∼ 18 ft and at the end, it becomes ∼ 22 ft. The pale red background color block shows the density of reinforcement learning strategy. This strategy is random because we have to look at the simulated (in reality, observed) sea level to decide each time step (as well as Bayesian learning which is not shown in the figure). We can find that reinforcement learning strategy start from a lower level than dynamic programming, which means this method is more "confident" about future risk control than traditional dynamic methods. The average final seawall height of the reinforcement learning result is on average ∼ 19.5 ft, which is between static optimal and dynamic programming result.

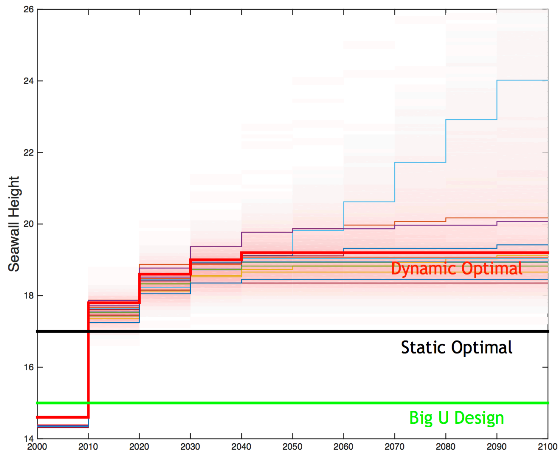

Figure 2: The Simulated Different Paths under different strategies

We can find the risk differences between different sea wall design strategies from Fig 3. The five curves presents the total cost of damage and construction of "Big U" design, reinforcement learning, Bayesian learning, dynamic programming and static optimal in legend order. The chance for New York City to suffer a 1 billion dollar damage (the orange vertical line) in the future is ∼ 10% for static and dynamic programming methods while the Bayesian learning and reinforcement learning methods show only about 10^{-3} probability. There are still chance for static and dynamic programming methods to exceed the 20 billion total cost threshold while almost no chance for that to happen under Bayesian learning and reinforcement learning method. Given the same survival level , for example 10^{-4}, we can find the quantile damage for Big U/ static optimal/dynamic programming are 12/4/2.5 times over reinforcement learning, which shows the excellent risk control ability of reinforcement learning.

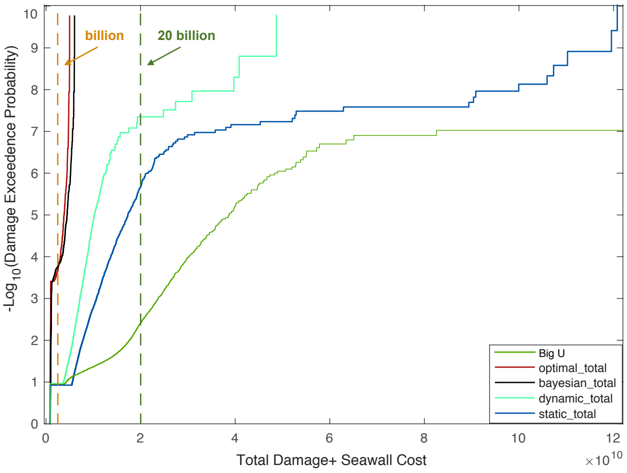

Figure 3: The Tail Part Risk Distribution for Life-cycle Economic Damage simulated under different Strategies

The expected total costs for the twenty-first century also different from each other. NYC will suffer 1.75 billion dollar damage under Big U design, while 1.37 for static optimal, 1.12 for dynamic programming, 1.05 for Bayesian learning and 1.01 for reinforcement learning. From this, we know that under the deep uncertain climate changing the environment, a slight change in our coastal protection design will lead to a 30% progress (from 15 ft Big U to 17 ft Static Optimal). Moreover,we quantified that economic-based adaptive design can earn 20% (comparing dynamic programming with static optimal) to 40%(comparing reinforcement learning to static optimal) for the NYC coastal protection problem. If the adaptive concept can be applied to all the climate change threatened infrastructure systems (even roughly), it can help people to save a lot.

These results show significant improvements from the current design. However, the improvement will be more considerable after coupling of current hurricane information. In this paper, the sizeable nonlinear hurricane transition is not included because for current technique the hurricane transition is hard to observe. In future condition, maybe we can observe that and thus will improve the ability of our reinforcement model.

5. CONCLUSIONS

In this article, we carefully analyzed the optimization model for adaptive seawall design and applied that to NYC. Based on the expected life-cycle cost, it shows a strict rank between reinforcement learning method, Bayesian learning, dynamic programming and static optimal. It is because of different information and different probability measure different methods have and also because different extents for them to use the observed available information. And also, we show that adaptive methods have several qualitative advantages. First, the cost distribution for learning based adaptive methods are usually lower than other methods; this is because the adaptive methods can respond to the observed climate information and adjust their strategies, which helps them to control the global risk. Secondly, we can find that the initial seawall height suggestions for adaptive methods are lower than static methods. This is because for static methods, once their strategies are determined, they have to cover the deep risk for a century or more, which leads them to a conservative position. However, an adaptive method can adjust themselves in the future, so they are more "aggressive". As both the deep uncertainty and the high cost are two main concerns for the government to start climate risk mitigation projects, the adaptive design seems to point out a new way for the government to tune the risk. For a government, they may be more willing to organize a policy or project under the lower risk sponsored by adaptive design. And also, with a lower initial cost, governments may be more willing and easier to gather resources and have a try on the projects. This also may simulate inter-country cooperation on climate agreements, because it can lower the climate protection requirements for member countries than the requirements under a static evaluation of policy nowadays.

6. ACKNOWLEDGEMENT

This material is based upon work supported by the National Science Foundation (grant 1652448). We also thank Prof. Warren Powell for his help on explaining BADP and Prof. Robert E. Kopp for future sea level simulations he provided.

7. REFERENCES

Adam, B. and Smith, I. F. (2008). "Active tensegrity: A control framework for an adaptive civil-engineering structure." *Computers & Structures*, 86(23-24), 2215–2223.

Aerts, J. C., Botzen, W. W., Emanuel, K., Lin, N., de Moel, H., and Michel-Kerjan, E. O. (2014). "Evaluating flood resilience strategies for coastal megacities." *Science*, 344(6183), 473–475.

Emanuel, K. A. (2013). "Downscaling cmip5 climate models shows increased tropical cyclone activity over the 21st century." *Proceedings of the National Academy of Sciences*, 110(30), 12219–12224.

Haasnoot, M., Kwakkel, J. H., Walker, W. E., and ter Maat, J. (2013). "Dynamic adaptive policy pathways: A method for crafting robust decisions for a deeply uncertain world." *Global environmental change*, 23(2), 485–498.

Kopp, R. E., Horton, R. M., Little, C. M., Mitrovica, J. X., Oppenheimer, M., Rasmussen, D., Strauss, B. H., and Tebaldi, C. (2014). "Probabilistic 21st and 22nd century sea-level projections at a global network of tide-gauge sites." *Earth's Future*, 2(8), 383–406.

Lickley, M. J., Lin, N., and Jacoby, H. D. (2014). "Analysis of coastal protection under rising flood risk." *Climate Risk Management*, 6, 18 – 26.

Lin, N., Emanuel, K., Oppenheimer, M., and Vanmarcke, E. (2012). "Physically based assessment of hurricane surge threat under climate change." *Nature Climate Change*, 2(6), 462.

Lin, N., Kopp, R. E., Horton, B. P., and Donnelly, J. P. (2016). "Hurricane sandy's flood frequency increasing from year 1800 to 2100." *Proceedings of the National Academy of Sciences*, 113(43), 12071–12075.

Lin, N. and Shullman, E. (2017). "Dealing with hurricane surge flooding in a changing environment: part i. risk assessment considering storm climatology change, sea level rise, and coastal development." *Stochastic Environmental Research and Risk Assessment*, 31(9), 2379–2400.

NYCGovernment (2017). "Nyc: The big u | rebuild by design.

Powell, W. B. (2007). *Approximate Dynamic Programming: Solving the curses of dimensionality*, Vol. 703. John Wiley & Sons.

Ranger, N., Reeder, T., and Lowe, J. (2013). "Addressing 'deep'uncertainty over long-term climate in major infrastructure projects: four innovations of the thames estuary 2100 project." *EURO Journal on Decision Processes*, 1(3-4), 233–262.

Rasmussen, D., Meinshausen, M., and Kopp, R. E. (2016). "Probability-weighted ensembles of us county-level climate projections for climate risk analysis." *Journal of Applied Meteorology and Climatology*, 55(10), 2301–2322.

Shultz, J. M. and Galea, S. (2017). "Preparing for the next harvey, irma, or maria—addressing research gaps." *New England Journal of Medicine*, 377(19), 1804–1806.

van der Pol, T. D., van Ierland, E. C., and Gabbert, S. (2017). "Economic analysis of adaptive strategies for flood risk management under climate change." *Mitigation and Adaptation Strategies for Global Change*, 22(2), 267–285.

Xian, S., Feng, K., Lin, N., Marsooli, R., Chavas, D., Chen, J., and Hatzikyriakou, A. (2018). "Brief communication: Rapid assessment of damaged residential buildings in the florida keys after hurricane irma." *Natural Hazards and Earth System Sciences*, 18(7), 2041–2045.