SOLVING PANDORA'S MATRIX: A SEARCH THEORY MODEL*

Anirudh Viswanathan
University of California, Berkeley
December 2018
Senior Honors Thesis

Abstract

The purpose of this paper is to outline the development and eventual solution to a new type of search model, called Pandora's Matrix. The model allows the agent to search through boxes containing unknown prizes, with a T number of periods to claim these prizes. After the T periods expire, the agent receives the total sum of all of his or her collected prizes. The rest of the paper focuses on developing an optimal stop and search rule, for the risk-averse agent, and highlights points of further work and experimentation.

*** My appreciation for my research mentor, Shachar Kariv, runs deeper than I could ever put into words. Shachar has been a wonderful mentor, insightful professor, and understanding friend through everything the past two years.**

I. Introduction

Search theory has been a highly consequential and significant field of economics that grants economists insight as to how consumers and sellers interact during trades, and the various alternatives they seek in order to optimize their economic benefits. A most prominent example is the analysis of the job search of workers during times of unemployment. For the workers, they desire job opportunities that offer fair wages, external benefits, and safety in the work place. How these workers go about making these decisions and what they ultimately choose is the main motivation of the research in this field. One paper in particular, "Optimal Search for the Best Alternative" by Martin Weitzman introduces a general, yet powerful game called "Pandora's Problem." The main objective of the model is to create an optimal search and stop rule such that the agent can choose a box amongst a selection of boxes, each of which contains an unknown prize. The main intrigue of this model is for how long should their search continue, and when is the ideal time to stop search? Weitzman offers his solution, and its implications are quite large. Based on this idea, I developed my own search game based on a matrix setup where the player can sample different boxes, but is restricted to search only an arbitrary number of boxes out of the total offered, and his prize is the aggregate sum of all the prizes in the searched boxes. This model allows me to solve for a closed-form solution that conveys what the optimal value is for every possible search restriction. This value is what I deem a threshold value and is the optimal number that the agent must search for in order to then terminate search and take that prize.

As such, the following sections of this paper will be as follows: relevant literature outlining the basics of search theory, a detailed summary of two papers that were instrumental in the creation and solution of the model I developed, a proof for the closed-form solution to my new search theory problem, and finally, concluding remarks and ideas for further research and implementation.

II. Literature Review

Economists (especially microeconomists) have been fascinated with the idea of search behavior, as it grants insight into the decision-making processes of consumers, laborers, businesses, and more. Stigler (1961) was the first to publish on the idea of search theory, and proposed studying individuals' search for bargains/wages as an economic problem. Subsequent to Stigler's seminal work, John McCall analyzed, through the process of optimal stopping, which job offer an unemployed worker should take, and defined his optimal stopping rule based on the idea of a reservation wage. The reservation wage is defined as the lowest wage that the agent would be willing to work at, and McCall proposes that the worker should only accept a job offer if the wage is higher than his reservation wage. There are further nuances that other economists explored once the restrictions of McCall's model were relaxed, as we notice that reservation wage can fluctuate (for example, if the worker has remained unemployed for a long duration, their reservation wage might become lower due to desperation or stigma from outsiders).

Martin Weitzman (1979), in his highly influential paper, "Optimal Search for the Best Alternative" derives an optimal search and optimal stopping rule for a broad collection of search problems, that he deems "Pandora's Problem." The model and resulting rules that Weitzman presents are a large influence in the search problem that this paper develops and focuses on. Weitzman begins his paper by outlining an example of a research team trying to optimize production of their current technology, through either an alpha or an omega process. The alpha process has the higher expected value for company savings, but, somewhat counterintuitively, Weitzman proves that developing the omega process is the optimal policy. He claims "there is a crucial difference between the value of a project and the order in which it should be researched." After solving this example, "Pandora's Problem" is detailed for the reader. There will be n closed boxes, and the specific box, i, will have reward x_i given by probability function $F_i(x_i)$. Each box has a cost, c_i that is affiliated with opening it and the prize of the box is given with a time lag, t_i. The main purpose of the setup is that Pandora has to decide at each stage whether or not to open a new box, or collect one of the prizes that she has already attained. Pandora's main objective is to maximize her expected present discount value (in simpler terms, this is the highest reward value, x_i, factoring in the difference in value between present and future).

In reference to previous work, Weitzman addresses that when all boxes were identical, the optimal policy for searching and stopping was easy: search until you have found a prize higher than your reservation value. The search and stop rule with alternate search capabilities is analogous to that rule. Assume that there are only two boxes, a closed box, *i*, and the reward of the initial box, z_i. Either Pandora can collect z_i or, a net benefit of:

$$-c_i + \beta_i \left[z_i \int_{-\infty}^{z_i} dF_i(x_i) + \int_{z_i}^{\infty} x_i \, dF_i(x_i) \right].$$

In the special case of indifference, where the closed and open box prizes are the same value, we set z_i equal to the preceding equation and attain:

$$c_i = \beta_i \int_{z_i}^{\infty} (x_i - z_i) \, dF_i(x_i) - (1 - \beta_i) z_i.$$

The beauty of this equation, as Weitzman argues, is that, if we let z_i represent the expected present discount value of following an optimal search rule, then z_i will be the reservation prize of box *i* and "all relevant information about box *i* is summarized by z_i". Thus, Weitzman presents his selection and stopping rule:

"SELECTON RULE: If a box is to opened, it should be that closed box with the highest reservation prize.

STOPPING RULE: Terminate search whenever the maximum sampled reward exceeds the reservation prize of every closed box."

While seemingly simple, the prior statements are extraordinarily powerful, as everything about Pandora's Problem can be solved using the value of reservation prizes (which are calculated from the second equation). Another important conclusion from this is, when sampling, it is optimal to first sample from boxes with high variance (riskier) to strike rich quickly and terminate search early, and then move on to low-risk, high reward distributions. Obviously, and as Weitzman points out himself, there are certain limitations, given the stringency of his assumptions, and much room for further work. Yet, the results that he arrived at proved invaluable for the field, and spurred on further research.

Laura Doval's doctoral thesis offers great insight into the ramifications of Weitzman's model once we relax certain assumptions. The focus of her paper is predicated on the idea that Weitzman's assumption of only being able to select a box after searching its contents is unnatural. Instead, Doval proposes that the elementary Search and Stopping Rules are too simple, but can become more powerful if the prior assumption is relaxed. Doval gives an accurate example of a recently-admitted student who is choosing among three colleges to attend (or not to attend a college), and treats each college as a separate box. Obviously, the student is able to attend a college even if he has not visited (or opened the box). Doval, thus, tackles this problem through the same general setup as Weitzman, and adds a utility function $u(z) = \bar{z}$, where $\bar{z}$ is the highest coordinate in the vector, z, of realized prizes.

Then, the author delineates the function that determines whether or not the agent continues to search or stop, by $\varphi(U, z) \in [0,1]$, where if the function is 0, search is terminated, and if 1, search continues, and U is the set of uninspected boxes. Then, we define $\sigma(U,z)$ to be the box that is inspected next. Thus, upon inspection, the agent is at decision node:

$(U \setminus \{\sigma\}, z \circ x_\sigma)$, and then would select: $\varphi(U \setminus \{\sigma\}, z \circ x_\sigma)$ and $\sigma(U \setminus \{\sigma\}, z \circ x_\sigma)$. So, the optimal strategy, according to Doval, solves the following problem:

$$V^*(\mathcal{U}, z) = \max\{\bar{z}, \max_{i \in \mathcal{U}} \mu_i, \max_{i \in \mathcal{U}} -k_i + \int V^*(\mathcal{U} \setminus \{i\}, z \circ x_i) dF_i(x_i)\}.$$

where V is the payoff using the optimal strategy. After this introduction, Doval goes on to analyze the solution for the case of a single box.

In this instance, she considers two alternatives, when $\bar{z} \geq \mu_i$ and when $\bar{z} \leq \mu_i$. For the first relation, either the agent terminates search (and takes $\bar{z}$), or they decide to open the box *i if:*

$$\bar{z} \leqslant -k_i + \int_{-\infty}^{+\infty} \max\{x_i, \bar{z}\} dF_i(x_i) \Leftrightarrow k_i \leqslant \int_{\bar{z}}^{+\infty} (x_i - \bar{z}) dF_i(x_i).$$

To further elucidate this, Doval defines the reservation value of the box to be, x_i^R such that:

$$k_i = \int_{x_i^R}^{+\infty} (x_i - x_i^R) dF_i(x_i),$$

When the value of k_i is substituted into the prior equation, we get:

$$-k_i + \int_{-\infty}^{\overline{z}} \overline{z} dF_i(x_i) + \int_{\overline{z}}^{+\infty} x_i dF_i(x_i) = \int_{-\infty}^{\overline{z}} \overline{z} dF_i(x_i) + \int_{\overline{z}}^{x_i^R} x_i dF_i(x_i) + \int_{x_i^R}^{+\infty} x_i^R dF_i(x_i).$$

This shows that "reservation value represents the highest prize that the agent expects to obtain from inspecting…agent's payoff from inspecting box i is bounded above." Now, in the case of $\overline{z} \leq \mu_i$, the agent, if he or she chooses to stop, will take box *i* without inspection. In contrast to the Weitzman paper, a new term, deemed the "backup value" is defined as follow:

$$k_i = \int_{-\infty}^{x_i^B} (x_i^B - x_i) dF_i(x_i)$$

The importance of the backup value is that it represents the value of an outside option such that the agent is indifferent between inspecting the new box, *i*, and taking the box without inspection. Conceptually, this value is similar to the reservation value, and proves to be invaluable when answering the caveat added to Doval's model. In summary,

(*)

1. *If $\overline{z} \leqslant x_i^B$, the agent takes box i without inspection.*
2. *If $x_i^B < \overline{z} < x_i^R$, the agent inspects box i and takes the larger prize between $\overline{z}$ and the sampled prize, x_i.*
3. *If $x_i^R \leqslant \overline{z}$, the agent does not inspect box i and takes his outside option.*

Before presenting her final optimal search and stop rule, Doval gives three propositions that are used in the proofs of her optimal rules. Firstly, Proposition 1 states that if the backup value at decision node (*U,* z) is less than $\overline{z}$, then, we follow the course of Weitzman's optimal search and stop rules (however, this condition is not necessary for Weitzman's stopping rule to be correct). Then, Proposition 2 dictates that assuming $\overline{z} < x_i^B$ for some unopened box, *i*, then for if it is optimal to terminate search and take some other box *m* that remains unopened, box *m* has the highest mean, reservation, and backup value amongst the set of all unopened boxes. Lastly, Proposition 3 details that at decision node (*U, z*), if *l* is an element of U with highest reservation value, then there exists an element *j* in U such that $x_j^R < x_l^R$. If the max$\{x_j, \overline{z}\} \leq x_l^R$ then, it is optimal to inspect box *l* at decision node ($U \setminus \{j\}$, z ○ x_j), then it is not optimal to inspect box *j*. Intuitively, this displays an instance where the agent does not want to open the box with the highest reservation value (for example, as a means of a backup prize). If this is not the case, and the agent does indeed want to open one more box, it should be the box with highest reservation value.

After establishing all of this, Doval proceeds to split up the search order and stopping rule into a further two cases: (1) when the optimal policy is the same as Weitzman's for all but the last box, and (2) when the boxes have binary prizes and equal inspection costs. The second case is where the optimal policy has the most deviation from Weitzman's.

For (1), for any boxes, *i, j*, such that $x_j^R \leq x_i^R$, let $\prod_{ij}$ be the payoff for inspecting *i* and then applying (*) to *j* and $\prod_{ji}$ be the payoff for inspecting *j* and then applying (*) to *i*. Thus, Doval writes her search order rule to be "If a box is to be inspected next, it should be the box with the highest reservation value." The rule, for this scenario, is "If there is more than one box remaining, stop only if the maximum sampled prize is higher than the highest reservation value amongst uninspected boxes, and take the maximum sampled prize," followed by, "If only one box remains, stop if the maximum sampled prize is higher than x^R or lower than x^B. In the first case, take the maximum sampled prize; otherwise, take the remaining box without inspection."

Binary prizes is the last problem that Doval explains, and states that for every box, $X_i = \{y, x_i\}$, where *y* is strictly less than x_i and $p_i = P(X_i = x_i)$. For the analysis of binary prizes, the optimal policy cannot be solely calculated by juxtaposing the backup and reservation prizes as done in (1). The optimal policy is as follows: For n ≥ 1, say boxes {1, ..., n -1} have been inspected, and let $\overline{z}$ denote the maximum sampled prize. Then, the order in which the boxes are inspected is dictated by:

$$p_n x_n^R + (1 - p_n) v_n \geqslant W(\{n, ..., N\}_D),$$

Here, the agent inspects the box *n* if this conditions holds or otherwise inspects box *n+1*. v_n is the payoff of the optimal policy at decision node ({n+1,…, N}). W({*n,…,N*}) is the payoff obtained by following Weitzman's rule with an outside option of μ_n. Stopping occurs when $\bar{z} \geq x_n^R$ or, $\bar{z} < x_n^B = \max_{i \geq n}(x_i^B)$ and $x_n^B = v_n$, (where the agent will take box n without inspection). In short, Doval has given solutions to a reworked model of Weitzman, and her solutions are very noteworthy for the field.

Thus, I have summarized the origins of search theory as well as rigorously gone through two papers that served as a large influence to the problem I outline and solve in this paper. The next section will cover the details of the model I constructed and its implications.

III. Methodology and Design

As I had outlined in the previous section, the famous "Pandora's Problem" created by Martin Weitzman serves as a blueprint for much further research in search theory, and was the main muse for the creation of the problem in this paper. "Pandora's Problem" deals with an unknown probability distribution with a random number of boxes, where she has the ability to search through as many boxes as she chooses, and can terminate search once she was found her optimal box. In conjunction with this idea, my version of the Pandora Problem, called Pandora's Matrix, involves a matrix with m x n dimensions, where each element (x_{mn}) in the matrix is a random number under a independently and identically distributed (IID) discrete uniform distribution from 1 to 100, represented by $U(1,100)$. Each x_{mn} is hidden, until the agent has "clicked" on that element (analogous to Pandora opening a box). Obviously, when the individual element is revealed, it will be some number between 1 and 100, with each number having a $\frac{1}{100}$ chance in appearing. Thus, the expected value of each x_{mn} is then 50.5.

I introduce the idea of a time restriction through the variable T, wherein the agent only has a T number of clicks to explore/unearth elements in the matrix. Additionally, one can always click on a previous value again, as many times as they would like, at any point along the T-n time periods. As a general rule, we assume that the product, $m \times n$ is always larger than the value of T, the reason being that we cannot permit the player to have opened every box and then make a decision. Once T has expired, the agent receives the total sum of all the values that he or she has clicked on or selected. The purpose of this game, is for the player to maximize their total payoff. In particular, given a certain *T*, there exists an optimal t_x such that the agent will always choose to click on that value till termination of the game. In this particular model, we assume k_i, the inspection cost, is 0, for the sake of simplicity.

Figure 1 shows a rudimentary depiction of the beginning of the game, before the game has started. In this example, the agent has 7 clicks.

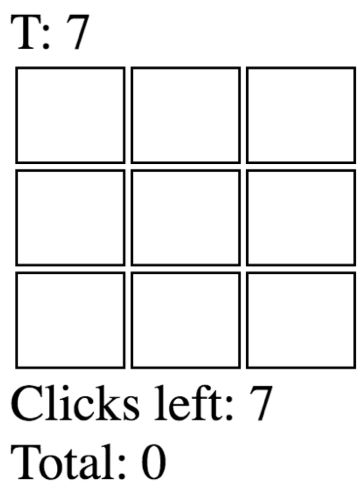

Figure 1: Beginning of game

Next, in *Figure 2*, the middle element, x_{22}, has been selected and gives the player a payoff of 77, with six clicks left. As you can see, the total is now 77, and with subsequent clicks, it will sum up all the numbers chosen.

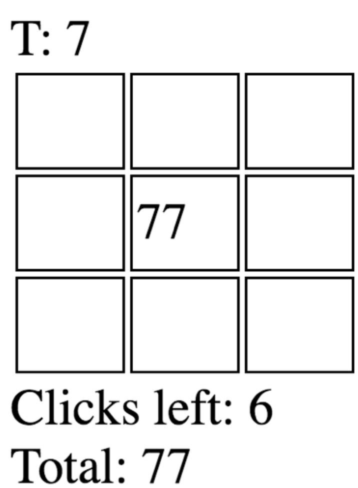

Figure 2: T-1 clicks left with initial payoff of 77

At this decision node, represented as *D(T-n, p)*, where *p* is the maximum of all values shown, the agent can either choose to click on 77 and receive a payoff of 154 at *T-2* periods (in this case, 5), or click on an open element and gain 77 + x_{mn} at *T-2* periods. As such, the game proceeds till termination, and the player ends with their total summation. Once again, at every *D(T-n, p)*, Pandora can either explore the matrix for a new value, or click on a previously shown value.

Figure 3 illustrates this, where at *D(6, 77)*, the agent decided to search for another value, clicks on x_{13}, and gets a value of 45. Unsatisfied with that, and with only 4 clicks remaining, 77 is taken till termination of the game, with a total payoff of 507.

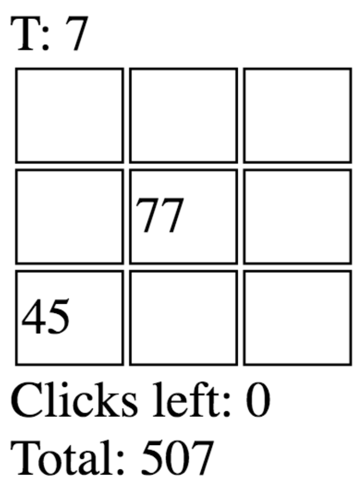

Figure 3: Termination of game after exploring matrix

Now that I have explained the main framework of the game, it is important to detail the main research questions that it leads to. Most importantly, as with the work of Weitzman and Doval, we want to develop our own stop and search rule. In this case, because of the restriction of the number of periods we have, the stop rule is already semi-enforced on the player. The player has an upward restriction of *T* to maximize their payoff. This leads to a contrast with the original Pandora's Problem, as is in this game, the agent must alter their search strategy contingent upon the number of clicks they have.

Thus, in light of the *T* click restriction, the idea of a threshold value is created. The threshold value is defined as the value such that once that value is attained (or a number above it), the player will stick with that value till termination of the game. Essentially, the threshold value serves as a minimum value that is inextricably tied with the number of clicks that are left. A precise mathematical definition of this will follow in the next section along with its connection with *T* and *p*.

IV. Optimal Stopping Rule for Risk-Neutral Agent

In this section I will delineate the subtleties of solving this new Pandora's Matrix problem, as well as go through a simple example of what a threshold value is, an explanation and mathematical proof for finding the threshold value for an arbitrary *T*, and conclude with a graph to show its increasing, asymptotic nature.

As stated previously, the objective of this game is to maximize the total payoffs at the end of the game. At every *D(T-n, p)* the player can either click on a previously uncovered value, or search the matrix for a new value, and each unclicked box has an E[X] of 50.5. Intuitively, this means that every node, if the previous max of all known values, *p*, is less than 50 (since we are in the discrete case), then, it is optimal to continue search. Unfortunately, this fact also belies the true intricacy of solving for an optimal stop rule, as it would seem obvious that search should cease once our agent has clicked on a value > 50. While it seems logically sound, it raises the question, what if $T = 5$ or $T = 10$, or any other large number. It would seem remiss to not even attempt to search for a higher value than 50 if a player had more time still available. As such, we can extrapolate that the optimal value shall change according to the total number of clicks left. Here is where the importance of the threshold value is displayed. The threshold value (t_x) tells the minimum optimal value needed, such that the agent remains with that value till termination. Obviously, having a value higher than our threshold value is better, and receiving a value that is lower than it, means it is optimal to continue search. We can define it such that, every decision node subsequent to finding the threshold will be $D(T\text{-}n\text{-}k, p) = \max(t_x, p)$. Thus, the total payoff will be some previous sum, $r + \max(t_x, p) \text{ X } x$, if x is the number of clicks left.

To elucidate this further, and the key idea behind the proof for the closed-form solution, let us consider the case where *T=2* and *T=3*. We will attempt to solve these through backwards induction. In the case of *T=2*, the agent has two clicks. They can either click on a x_{mn} that is greater than or equal to 50.5, in which case, the E[x_{mn} > 50.5] = 75.25. Obviously, with *T=1*, the agent will click on 75, again. In the opposite case, where x_{mn} is less than 50.5, E[x_{mn} < 50.5] = 25.25. With *T=1*, the agent will choose to search again, and click on a new element, with expectation of 50.5. In total, the agent has a 50% chance of having a total payoff of 75.75 or a 50% chance of having a total payoff of 150.5. Finding the average of that gives us 56.2525, which is $\frac{225}{4}$. Thus, for *T=2*, t_x = 56.25. Now, by using backwards induction, every game larger than *T=2* will always end up in the same scenario. So, anytime a player arrives at the time period of *T=2*, they know that unless they have already attained value greater than 56.25, they must continue search. For the case of *T=3*, the player will go through a similar progression, except now, the initial value that they will try and click on will be either greater than or equal to, or less than $\frac{225}{4}$. This is because we know that this is the value that must be attained in order to maximize the total payoff.

Given the examples above, I had to find a way to conflate the definition of what the threshold value was, along with the intuition of solving it as I would any sequential game. Ultimately, I realized that the way to find my closed-form threshold equation would be to analyze the game through the expected value with x turns left, in conjunction with whatever p is at that time. I denote this as $E_{x,p}$, which is the total expectation with x turns left. $E_{x,p}$ is thus equal to max(t_x, p) x x. Then, the proof is based on the following idea: if I click on a threshold value, there is an indifference case, where I can either click on the threshold value till termination, or find a new value (and then select the maximum between this new value, y, or the prior t_x).

The proof is as follows:

$$t_x \cdot x \geq 50.5 + \frac{1}{100} \cdot \sum_{y=1}^{100} E_{x-1,max(t_x,y)}$$

$$t_x \cdot x \geq 50.5 + \frac{x-1}{100} \left(\sum_{y=1}^{t_x} t_x + \sum_{y=t_x+1}^{100} y \right)$$

$$t_x \cdot x \geq 50.5 + \frac{x-1}{100} \left(t_x^2 + \sum_{y=1}^{100} y - \sum_{y=1}^{t_x} y \right)$$

$$t_x \cdot x \geq 50.5 + \frac{x-1}{100}\left(t_x^2 + \sum_{y=1}^{100} y - \sum_{y=1}^{t_x} y\right)$$

$$t_x \cdot x \geq 50.5 + \frac{x-1}{100}\left(t_x^2 + \sum_{y=1}^{100} y - \sum_{y=1}^{t_x} y\right)$$

$$t_x \cdot x \geq 50.5 + \frac{x-1}{100}\left(t_x^2 + \frac{100(101)}{2} - \frac{t_x(t_x+1)}{2}\right)$$

$$t_x \cdot x \geq 50.5 + \frac{x-1}{100}\left(\frac{t_x^2}{2} + \frac{100(101)}{2} - \frac{t_x}{2}\right)$$

$$t_x \cdot x \geq 50.5 + (x-1)\frac{101}{2} + \frac{x-1}{100}\left(\frac{t_x^2}{2} - \frac{t_x}{2}\right)$$

$$0 \geq \left(\frac{x-1}{200}\right) t_x^2 - \left(\frac{x-1}{200} + x\right) t_x + 50.5x$$

$$t_x \geq \frac{\left(\frac{x-1}{200} + x\right) \pm \sqrt{\left(\frac{x-1}{200} + x\right)^2 - \left(4 \cdot 50.5x \cdot \frac{x-1}{200}\right)}}{\frac{x-1}{100}}$$

$$t_x \geq \frac{200 \cdot \left(\frac{x-1}{200} + x\right) \pm \sqrt{200^2 \cdot \left(\frac{x-1}{200} + x\right)^2 - 200^2 \cdot \left(4 \cdot 50.5x \cdot \frac{x-1}{200}\right)}}{200 \cdot \frac{x-1}{100}}$$

$$t_x \geq \frac{(201x-1) \pm \sqrt{(201x-1)^2 - (404 \cdot 100x \cdot (x-1))}}{2x-2}$$

$$t_x \geq \frac{(201x-1) \pm \sqrt{40401x^2 - 402x + 1 - 40400x^2 + 40400x}}{2x-2}$$

$$t_x \geq \frac{(201x-1) \pm \sqrt{x^2 + 39998x + 1}}{2x-2}$$

$$t_x \geq \frac{(201x-1) - \sqrt{x^2 + 39998x + 1}}{2x-2}$$

The first line in the proof was set up based on the previous definition I had given. Then, I rewrote the sum of the expected value based on the probability of the new and old values, and altered the bounds on the summations to fit the new parameters. Through algebraic manipulation and Euler's summation formula, I was able to simplify the summations into a quadratic equation, which I then solved using the quadratic formula. It is important to note that the closed-form solution is undefined for the case of $x = 1$, because the basis of the proof is the ability to click on a new value. That is easily remedied by the fact that I can define a piecewise function, such that for $x = 1$, $t_x = 50.5$, which is obvious, as the result of the expected value of uniform distributions.

Figure 4 and 5 show graphs of this closed-form solution, and in particular, how there is an asymptote at the value of 100.

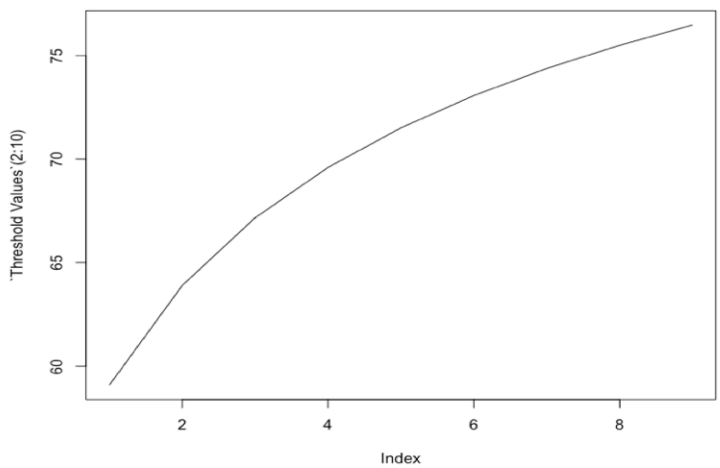

Figure 4: Graph of threshold values from T=2 to T=10

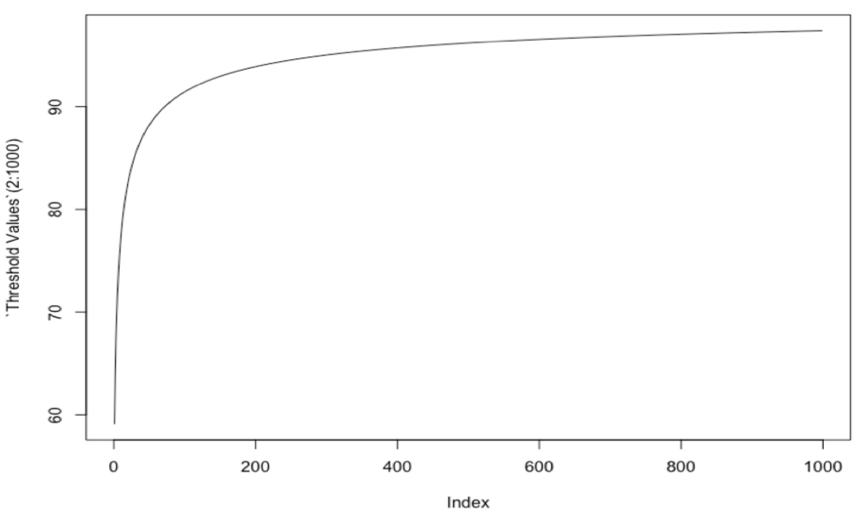

Figure 5: Graph of threshold values from T=2 to T=1000

V. Further Work and Concluding Remarks

The main purpose of this paper was to create and solve a unique variant of the age old problem developed by Martin Weitzman. I constructed a model with certain restrictions from the Pandora Problem, and added the concept of not only aggregating the payoffs, but also of limiting the total search time, with a period restriction of *T*. A uniform distribution of *U(1,100)* was used for the contents of each matrix element. This, in turn, introduced the idea of a threshold value which serves as the minimum value needed for optimal stopping and claiming. Through the definition of the threshold value, it was discovered that solving for the indifference case, wherein you can either continue searching or stay with the threshold value, gives a nice closed form solution. This closed form solution gives a full list of threshold values for any arbitrary number of clicks, and, as we can expect, it reaches an asymptote at 100, since intuitively, with a sufficiently sized matrix and a large number of clicks, an agent should search till they reach 100.

The model that I have constructed has many implications, and there is much further work to be done to use it to its fullest capacity. Namely, solving for the case of an agent that is risk-averse will add a new wrinkle into the problem, as now must consider concave utility functions, as opposed to the constant utility a risk-neutral agent feels. The power of Pandora's Matrix Problem is also its capacity to be tested on a group of humans and see how rational they are in making these optimal stop and search choices. This can, in turn, give us insight into the decision making of various people, and offer an empirical foundation for this problem, as well. Furthermore, a real-world example of this problem could be if a company is searching for a group to outsource their labor or construction, and after searching through different groups, they decide upon one for every subsequent project. Another interesting idea would be a situation where the agent has the previous options that they selected for only a set amount of turns before they disappear. Meaning, if the agent is unhappy with everything they have clicked previously and continues searching, they lose the ability to go back and use what they have already seen before. This would force the players to be more decisive as well as change the idea of the threshold value, as the set number of clicks is then broken into smaller games. An immediate corollary to this would be setting a time lag of the matrix elements, such that the agent only discovers the values that they have attained a few turns after they have clicked on it. All of these different iterations of the game have great economic implications, which will only aid in the subsequent studies and insights given by search theory.

VI. References

George Stigler. The economics of information. *Journal of Political Economy*, 69:-225, 1961.

H. Chade and L. Smith. "Simultaneous search". *Econometrica*, 74(5):1293–1307, 2006.

J. J. McCall *The Quarterly Journal of Economics*. Vol. 84, No. 1 pp. 113-126, 1970.

Laura Doval. "Whether or not to open Pandora's box," Journal of Economic Theory, Elsevier, vol. 175(C), pages 127-158, 2018.

M. Weitzman. "Optimal search for the best alternative". *Econometrica: Journalof the Econometric Society*, pages 641–654, 1979.

S. Lippman and J. McCall. "The economics of uncertainty: selected topics and probabilistic methods." *Handbook of Mathematical Economics*, volume 2, chapter 6. North-Holland Press, 1982.

W. Olszewski and R. Weber. "A more general pandora rule?" *Journal of Economic Theory*, 160:429–437, 2015.