Optimal Stopping Under Stochastic Observation Conditions

David B. Menasche

American University Department of Economics

Adviser: Mary E. Hansen

Abstract

You have a series of options. You need to decide whether to keep looking for more choices, or to pick from among those you've already discovered, but you don't know what else is out there. What is the optimal time to stop looking and pick? This is the thrust behind this work, in which I develop a theory of optimal stopping consistent with conventional utility maximization and study how the distribution of options in the potential choice space affects this maximization. Combining a mathematical treatment with computer simulation reveals that the searching agent systematically fails to achieve theoretically possible maximum utility while following a rational decision making model. The magnitude of this failing is found to be linearly dependent on the variance in utility of the distribution of choice space options. What this means operationally is that having a more uniform spread in quality of options can reduce happiness; applications include job searching, home and asset sales, and dating behavior.

I.

Sometimes people don't like making decisions. It doesn't really matter what the actual decision is, anecdotally it's not uncommon for people to "go with the flow" and defer a choice to a friend, family member, or spouse. And there's an understandable reason. Imagine deciding where to take your family to dinner in an area you don't know well. Not only may you feel external pressure from your family members to make a good decision—based on price, tastiness of the food, atmosphere, etc.—you also have to take time to look at and sort through the available choices. True, your time is valuable, but the act of looking has other costs too, including gas and energy in this case. It's completely reasonable for the other parties involved to want to avoid these costs if they can. By implication, whatever disutility is associated with the task of searching through available options and coming to a decision outweighs the utility of actually being able to pick.

Given then that you're in charge of choosing the best option from an unknown set, you want to balance the benefits of looking against the costs. When the costs outweigh the benefits, it's in your best interest to stop. Finding that point is the process of optimal searching, and that point is the solution to the optimal searching problem.

In this paper, I look to develop a general theory of optimal searching under certain conditions, namely the following: 1) the searching agent has rough but imperfect knowledge of the distribution of the value of each option; 2) the agent acts rationally and is able to cardinally rank options he finds; 3) the values of the found options are derived from a normally distributed random variable; and 4) the agent may look through infinitely many options should he choose to do so.

This general theory will aim to answer a number of qualitative and quantitative inquires relevant to searching. Given the conditions listed above, whether there is an explicit solution for the optimal time spent searching is of particular interest. Agents may use this knowledge to maximize their utility when confronted with a search in a general context. Further, how do the particular characteristics of each agent affect the optimal time spent searching? Surely an individual's personality profile affects the process in some way. And what about the characteristics of the underlying distribution option values? How does varying the distribution parameters optimal stopping point? In this paper I will aim to answer these questions and clarify the dependencies among these variables.

II. Literature Review:

Economists have worked to develop specific solutions to the optimal stopping problem as it relates to certain canonical issues. Among these, labor market dynamics, utility maximization, and financial option exercise stand at the foreground, and these cases are treated with tools that generalize to more a general theory.

Modeling labor market dynamics has proven difficult in the past due to the relative failure of classical labor supplied-labor demanded models to predict market clearing wage rates and unemployment levels. Under realistic conditions, an agent receives one job offer at a time and must choose within short time whether to accept the terms of employment. This decision involves estimating the chance of receiving a better paying offer and weighing that chance against the disutility associated with further search costs. This process is developed by Pierre Cahuc and Andre Zylberberg (2004), who use marginal utilities and differential equations to create families of solutions the problem. Even then, these solutions don't deal well with additional real-world complexities.

Nobel Laureate Peter Diamond laid the groundwork to deal with these complexities in his early 1980's study of search friction; having quantified the effect of a relatively small amount of economic inefficiency, he found that that inefficiency led to a relatively large disturbance from market theoretical equilibrium (Peter Diamond 1982).[1] These forays were further advanced by co-laureates Dale Mortensen and Christopher Pissarides, and were ultimately combined into the DMP model of searching and matching. While these models advance the case of multiple agents searching simultaneously for like "goods," they do not deal with the specific microeconomic decision made by an individually searching agent (Harold Cole and Richard Rogerson, 1999).

To this end, there is much literature that deals with the factors that underlie utility maximization inherent in this decision, when an agent is subject to certain conditions. It turns out that the gathering of information about available options is an especially nontrivial step in the process. Before the agent considers any of the options available to him he must decide whether he needs more information to make an informed choice about whether to continue searching (Wolfgang Stadje, 1997), and if necessary, must sacrifice some marginal amount of utility to make this determination.

The assumptions on which a utility maximizing agent with incomplete information will base his actions (Barton L. Lipman, 1999) also play an important role in the searching process. While people have demonstrated that they will generally make consistent choices 75% of the

[1] This phenomenon is known as the Diamond Paradox.

time, Blavatskyy has shown that people will make random errors in assessing the expected utility of an option the remaining 25%, indicating that actually determining the value of a given option is a particular difficulty of searching agents (Pavlo R. Blavatskyy, 2007). Intuitively, errors in valuation have been shown to be heteroskedastic (Blavatskyy, 2009).

Besides the intrinsic value of the searched-for option, other factors affect an agent's utility during the search process. An agent's patience (John Quah and Strulovici Bruno, 2009) limits the allowable duration of the search by changing the functioning the functional form and magnitude of the parameters in the utility discounting function. (As the search time mounts, so do opportunity costs, hence time-disutility begins to play a role.) At the same time, searching through too few options also has potential for disutility. When evaluated for the case of agent with an infinite time horizon choosing when to exercise a financial option, the agent is under constant temptation to realize a profit, provided the option is in the money (Jianjun Miao and Junjian Miao, 2008). If he exercises the option before the option reaches maximum value he will have lost out.

Compounding the difficulties in determining the best time to stop the search are regret effects. In general, people lose utility when they realize they would have been better off doing something besides what they are doing (Dean P. Foster and Vohra Rakesh, 1999). While people do have a tendency to disregard and minimize the disutility of future regret (Francesco Drago and Kadar Dora, 2006), a rationally searching agent should continually calculate this effect, less he halt a search prematurely.

III. <u>Theoretical Model:</u>

One can view the optimal stopping problem as a specific case of utility maximization. In general, given certain constraints and parameters, an agent must execute an action at a time that maximizes welfare. Under this model, this executable action is the choice to stop searching.

A. Assumptions and Definitions

I start by assuming the agent is in a situation where he may take an arbitrarily long time to search through various utility giving options, should he choose to do so. To begin, assume that the agent takes a certain discrete period of time to discover one of these options (through exploring his "observation space," the collection of all possible options, denoted henceforth by Ω), and that he can assign immediately a cardinal utility score to this option. Denote these utility giving options, henceforth "observations," by x_i, $i \in \{1, 2, \ldots, n\}$, where n is the total number of discovered observations at time t, and define a map $E: x_i \mapsto \mathrm{R}$, such that $E[x_i] = U_i$, where U_i is the utility of observation x_i. We may collect the discovered observations in a set called the "consideration space" denoted ω, such that $\omega = \{x_1, x_2, \ldots, x_n\}$. Let the i^{th} element of ω be denoted by $\omega[i]$.

Assume further that the utilities given by these observations are distributed normally by $\sim N(\mu_{pop}, \sigma^2_{pop})$ and that the agent has only rough knowledge of the distribution parameters. The agent's prior knowledge is given by $\sim N(\mu_{prior}, \sigma^2_{prior})$. In the simplest case, once the agent discovers an observation, we allow him to return to it and choose it cost-free, should he choose. This condition may be relaxed after further exposition.

B. Introducing the Utility Function

Searching comes with inherent costs. Intuitively, the agent should continue to search through the observations until the costs of searching for new options outweigh the benefits. In one simple model, there are two components that have the potential to change the agent's utility, holding constant $E[x_i]$. One of these factors is time-disutility. The act of searching takes time and time is a valuable resource. Accordingly, there should be some disutility inherent to the investigation of new observations. The other is disutility associated with regret; if an agent is too cursory in his survey of the observation space he may experience disutility from wondering if he has made the right choice by stopping the search.

Given this intuition, define a utility function with parameters ω, n, and t, where utility, U, is given by

$$(1) \qquad U(\omega,n,t) = Max(E[\omega]) - a_1 t - \frac{a_2}{n},$$

where a_1, a_2 are positive constants and t, n are positive integers. One could view a_1 as an agent's impatience—at higher values, one time unit will discount more utility. In the same way, a_2 may be viewed as an agent's propensity to experience regret. At higher values, an agent feels the need to investigate more observations. It is reasonable to adjust the units of time such that the agent discovers one observation per unit t. Note that this assumes the agent discovers new observations at a constant rate. Doing so and allowing the agent to find the first observation at time $t = 1$, then $t = n$, and U becomes

$$(2) \qquad U(\omega,t) = Max(E[\omega]) - a_1 t - \frac{a_2}{t}.$$

C. Maximizing Utility

The agent should choose the observation from his consideration space that yields the highest utility, hence the $Max(E[\omega])$ term. Examining the limiting behavior of the function agrees with intuition. As $t \to \infty$, the agent examines all the options (assuming Ω has infinitely many elements) and the regret term disappears; however in doing so, $U \to -\infty$ as the costs of searching mount. Figure 1 depicts a plot of (2) as a function of time, holding ω constant. Qualitatively, we see that U reaches a maximum when the benefits of alleviating regret are outweighed by the time disutility. The quantitative intersection in this figure is unique to the parameters chosen to generate the graph and is not yet important. Because we hold ω constant, we can also differentiate U with respect to t and optimize the function to solve explicitly for the optimal stopping time. Solving, we find

$$\frac{dU}{dt} = -a_1 t + \frac{a_2}{t^2}, \tag{3}$$

$$t_{ideal} = \sqrt{\frac{a_2}{a_1}}, \tag{4}$$

and

$$U_{\max} = Max(E[\omega]) - 2\sqrt{a_1 a_2}. \tag{5}$$

This is a simple result, but it hinges on the assumption that the agent's consideration space grows with no variation in the utility of each discovered observation. In any meaningful search the consideration space grows with observations that yield utilities distributed in some way. As mentioned, we will assume here that these utilities are distributed normally with parameters μ_{pop} and σ^2_{pop}. Solutions given by (4) and (5) are accurate in the absence of any variance in the underlying distribution of the observations, in other words as $\sigma^2_{pop} \rightarrow 0$. In the empirical analysis, I will aim to demonstrate how introducing some randomness into the observation space lowers the agent's values for t_{ideal} and $U_{\max}$, and how relaxing certain assumptions further affects the searching process.

D. The Search Simulation

The searching process will be simulated within Mathematica; I will begin by describing each step of the program. Before the agent interacts with his environment at all, I make the assumption that he has some rough knowledge of the underlying distribution parameters μ_{pop} and σ^2_{pop}, and achieve this computationally by generating two random numbers within an arbitrary interval and multiplying each by μ_{pop} and σ^2_{pop} respectively to establish μ_{prior} and σ^2_{prior}. While

μ_{prior} is free to vary above or below μ_{pop}, σ^2_{prior} should intuitively be larger than σ^2_{pop} given the agent understands that his knowledge is potentially inaccurate. Widening this interval accordingly and graphing the results yields something like Figure 2, with $x_i \sim N(\mu_{pop}, \sigma^2_{pop})$ graphed in black and $x_i \sim N(\mu_{prior}, \sigma^2_{prior})$ graphed in red.

The agent then begins the searching process. Assuming the utility function of form (2) the agent discovers x_1 at time $t = 1$. The agent will continue to discover new observations when the utility associated with discovering the next observation (including search costs) is greater than the current utility, thus one must define a Utility function, U_{t+1}, for the next time period. Generally,

(6) $$U_{t+1}[\omega_{t+1}, t, n_{t+1}] = Max(E[\omega_{t+1}]) - a_1(t+1) - \frac{a_2}{n_{t+1}},$$

where ω_{t+1} is the consideration space at time $t + 1$ and n_{t+1} is the number of total observations at that time. Given that the expectation of a normal distribution is the mean of that distribution, at time $t + 1$ the agent anticipates discovering an observation x_{t+1}, such that

(7) $$E[x_{t+1}] = U_{t+1} = \overline{U}_{\omega_t} = \mu_t,$$

where $\overline{U}_{\omega_t}$ is the mean utility of the observations in ω at time t and μ_t is convenient notation. It follows from (6) that

(8) $$U_{t+1}[\omega_t, \mu_t, t, n_{t+1}] = Max(E[\omega_t + \{\mu_t\}]) - a_1(t+1) - \frac{a_2}{n_{t+1}},$$

but we can note that even in general, $Max(E[\omega_t])$ is already given, so the first term is known if the agent can solve for μ_t. For the case where observations are discovered at a constant rate and $t = n$, (8) simplifies to

(9) $$U_{t+1}[\omega_t, \mu_t, t] = Max(E[\omega_t + \{\mu_t\}]) - a_1(t+1) - \frac{a_2}{t+1},$$

and the primary theoretical concern becomes solving for μ_t. To begin, we consider the case where $t = 1$, and the agent has discovered only one observation. The agent has rough prior knowledge of the distribution of Ω, but cannot immediately calculate μ_t without incorporating the new knowledge gained from the observation he discovered. He must revise this prior knowledge to reflect actual data about the true distribution of Ω. It is intuitive that given only one data point the agent cannot immediately take $x_1 = \mu_t$; instead, μ_t should be formed by averaging the point in some way with the prior knowledge. To do this, one option is for the agent to weight the observations by Z-score. Intuitively, if $E[x_t]$ is farther away from μ_{prior}, the agent should revise his prior knowledge more because he perceives it is more likely that his initial knowledge is inaccurate. As the agent discovers more observations, however, the strength of the previously computed mean should increase sue to the amount of supporting data.

Difficulties arise, however, because in order to compute the Z-score, the agent must compute some measure of the standard deviation that includes information from both his prior knowledge and the new data point. The standard deviation, in turn, depends on the current distribution mean so the two quantities must be computed simultaneously. Noting that for the first time period $\mu_{t-1} = \mu_{prior}$, this would imply simultaneous equations of the form

(10) $$\mu_t = \frac{(n)(\mu_{t-1}) + (z)E[x_t]}{n + |z|}$$

and

(11) $$\sigma_t = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(\omega[i] - \mu_t)^2}$$

where

(12) $$z = \frac{x_t - \mu_t}{\sigma_t},$$

however, until 'enough' observations are discovered, the agent cannot rely on a measure of the standard deviation calculated in this way. Certainly, the standard deviation of only x_t is not perfectly representative of the underlying distribution, so the agent should use his prior knowledge to construct some kind of average. Ideally, such an average would converge to the calculated value of the standard deviation when the number of observations in the consideration space is sufficient. Constructing such a function yields a piecewise function given by

$$\sigma_t = \begin{cases} \frac{N_0 - n}{N_0}\sigma_{prior} + \frac{n}{N_0}\sqrt{\frac{1}{n}\sum_{i=1}^{n}(\omega[i]-\mu_t)^2}, & n \le N_0 \\ \sqrt{\frac{1}{n}\sum_{i=1}^{n}(\omega[i]-\mu_t)^2}, & n > N_0 \end{cases}, \tag{13}$$

where N_0 is a parameter characteristic to agent describing in some way his confidence in his prior knowledge or otherwise the strength of his memory. If N_0 is relatively large, he will have to discover comparatively many observations before he disregards his old conception of the distribution of Ω and uses the his calculated value exclusively. If $n \le N_0$ then from (13) we can see that then the prior and the calculated values are combined with weights consistent with the number of observations discovered. After this point, when $n > N_0$, the agent will rely only on the value he calculates from the consideration space.[2]

Solving (10), (12), and (13) simultaneously yields values for μ_t and σ_t that the agent then uses to form his revised distribution, the mean of which will serve as the basis for the expectation of the next observation's utility. To demonstrate in the context of the previous figure that shows the agent's prior generation, Figure 3 displays the acquisition of the first observation and the

[2] Figure 3 displays the results of this revision process after 25 iterations; the underlying distribution of Ω is graphed in black, the agent's initial guess at the distribution of Ω is graphed in blue, and his guess after 25 iterations is graphed in red. The lighter curves show the process of convergence to the red plot.

agent's revision of the distribution. The dashed vertical line is $E[x_1]$, the black curve is the distribution of Ω, and the red curve is the agent's previous conception of the distribution of Ω. The utility of the observation generated is close to μ_{pop}; notice μ_t is closer to μ_{pop} and the discrepancy of the observation with respect to the prior seems to be about one standard deviation. Accordingly, the deviation of the revised distribution is marginally less.

As a slight digression, it is reasonable to ask whether the revision process is an accurate one. Using tests of 25 iterations each, I have found that very near convergence occurs quickly, regardless the observations generated. Figure 4 depicts three such cases, with the underlying distribution curve drawn in bold black, the agent's prior knowledge in bold blue, and the revised distribution after 25 observations in bold red. The intermediary revisions are the lighter curves. For this figure, the agent's value of N_0 is 10.

Using the same coloring scheme, Figure 5 shows the same convergence, but with a change in parameter, namely $N_0 = 5$. Convergence is more variable, but still appears acceptable.

After the agent revises his prior knowledge, he computes the expected utility of discovering another observation using (9). If it is higher than his utility currently, he continues to look. For the observation and priors shown in Figure 3, this process is depicted in Figure 6. In this figure, the expected utility at time $t + 1$ is graphed in red at time t. Current utility is graphed in blue. Provided the red dots are above the blue dots, the search should continue. The point in time at which the dots cross is the time at which the search would stop. Figure 7 depicts the agent's conception of Ω (in magenta) at the time the search initialized from Figure 3 ultimately concludes. It has converged to near to the actual distribution of Ω. Figure 8 displays the corresponding graph of Utility vs. Time at the optimal stop time.

Generating more observations and examining the agent's utility response as time goes on shows that the agent would stop the search prematurely; utility continues to increase despite the agent's projections. Figure 9 shows this progression. A full printout of this simulation run-through complete with code is available in the Figures section as Figure 14. The next section will quantitatively examine the relationship between the variance of Ω and the difference in utility achieved by the agent given the premature stop.

IV. **Emperical Analysis:**

Consistent with (4), there exists an optimal time to stop the search given no variance in the utility of the observations discovered and added to ω. (In other words, the agent discovers an observation that yields a given utility again and again.) We might then ask how introducing variance to Ω affects this optimal time.

Rerunning the simulation for various values of σ^2_{pop} (25 iterations per value) and averaging the difference in time of the agent's stopping against the no variation value given by (4) yields results given in Table 1. When this time difference is graphed as a function of the σ_{pop} /μ_{pop} ratio, the result is given by the curve graphed in Figure 10. The dependence is clearly nonlinear, but qualitatively, one notes that higher levels of uncertainty in the underlying distribution shorten the time the agent spends searching. The graph asymptotes to a value of ($t_{\max}$ – 1) because when variability in Ω gets to a certain level, the search ends most frequently in only one round and the maximum the difference in time periods can be is, in fact, ($t_{\max}$ – 1).

Pursing this angle of analysis further leads one to more questions. While we have already that established that the search ends before the no-variance theoretical value of t, it is also apparent from the qualitative shape of the agent's realized and expected utility vs. time graph (Figure 9) that the simulated search that includes variance also ends before it should. The agent could actually improve his ultimate stopping utility if he searches past the point where his expected utility (graphed in red) crosses his currently realized utility (graphed in blue) to the point where the realized utility is at a maximum. How much does stopping the search rationally (where expected utility crosses realized utility) lower the ultimate stopping utility, given that the realized continues on to another maximum?

It turns out that the size of the gap between the true maximum utility and the rational stopping utility is a linear function of the σ_{pop} / μ_{pop} ratio. This relationship is displayed in Figure 11; data used in least-squares regression are summarized in Table 2. The resultant best-fit line is given by

$$|\Delta U| = (3.68 \pm 0.09)\left(\frac{\sigma_{pop}}{\mu_{pop}} * 100\right) - 0.01, \tag{14}$$

where $|\Delta U|$ is the magnitude of the reduction in utility in percent. (To generate these data, 100 iterations of the model were performed at each value of σ_{pop} / μ_{pop}, with the results averaged.) These findings imply that for every 1% change in the underlying distribution, Ω, there is approximately a 3.68% reduction in stopping utility as measured by the vertical difference between the expected utility/realized utility crossing point and the level at which realized utility is actually maximized. See Figure 12 for a clarification of this point.

Again, as a consequence of the agent acting rationally (and consequently stopping his search too early), he does loses out, to the tune of 3.8% per 1% in the σ_{pop} / μ_{pop} ratio. The agent should look longer; one possible extension to the model could be to find out by how much.

Strangely, this result also points toward the conclusion that by minimizing the variance in an underlying distribution of options, the agent can achieve greater happiness for a given μ_{pop}. This is not intuitive—somehow the idea that an agent might gut out a longer search and try to get lucky is hard to shake. It is important to note, however, that the values for time-disutility and regret parameters a_1 and a_2 are held fixed in these cases. It is possible that somehow if the agent expects a wider distribution (with extreme values more likely) his time-disutility coefficient will be lower indicating more patience and a willingness to look for longer. Another interesting extension would be to make the parameter a_1 randomly generated but inversely proportional to σ_{pop} and thus σ_{prior}. Without this addition, it suffices to say that more variability in Ω tricks the agent into stopping sooner to some extent.

An easy way to improve the applicability of the searching process is to inquire whether the qualitative effects of the previous differ in nature to a search where the agent cannot immediately return to an observation once he has moved on. This is analogous to an individual choosing whether to accept a job or sell a house given an offer; if the expected value of a later observation is higher once the costs of waiting and searching are subtracted, it is rational for the individual to wait. Intuitively, these searches should take longer. The agent develops knowledge of the distribution at the same rate in both cases, but it is more likely the agent in the second case will keep searching because half the time he will be holding a below-average observation that he would rather trade for another of better value.

Given the new restrictions in the searching process, one might then imagine a scenario where a rationally acting agent moves on from an observation expecting to find something better but is then disappointed to find the observation he relinquished was actually above average. It might then take him considerable time to recover an observation of the same value; given the agent is rationally acting and has some threshold value for the time-discounting parameter, he may be forced to stop before he finds another due to time disutility. This is a settling scenario:

the agent knows that better and more valuable observations are present in Ω but has looked long enough that he must make a less optimal selection.

Indeed, this circumstance comes to fruition. Depicted in Figure 13, the agent's realized utility is no longer monotonically increasing. At $t = 6$, the agent expects a higher utility but is disappointed and is forced to ultimately accept a stopping utility less than his utility at $t = 6$. To construct this example, iterative runs of the model were performed using different values of the time discounting parameter to determine the threshold where settling would occur. The scenario first manifested itself in the results when the a_1/μ_{pop} ratio was 0.15.

V. **Conclusion and Further Research**:

Due to the stochastic expansion of ω, it is not possible to solve for an explicit expression that describes the utility-maximizing time while there is variability present in Ω. That said, it is possible to iterate the constructed model numerically and note qualitative behavior and commonalities among iterations. Common to all trials is the observation that if the agent follows a rational heuristic—comparing next time period's utility to this period's—he will stop searching before he realizes his maximum possible utility. The amount of time he must continue to reach this point remains a viable point for further research.

Despite this shortcoming, the size of the gap between stopping utility and maximum utility has been quantified and found to be directly proportional to the variance of the underlying distribution of observations. This implies that an agent, given a choice between two observation spaces to search through, should pick the one with less variation given each has the same value for μ_{pop}. When parameters are exogenously given and the agent searches consistent with current restrictions, less choice may actually mean more utility, ultimately.

This may change if utility function parameters a_1 and a_2 are functions of the underlying distribution or functions of the priors generated from the underlying distribution. An intuitive extension is to create some dependence between a_2, the regret coefficient, and σ_{prior} or σ_t. If an agent believes there to be more variability in Ω and he may experience more regret for a given number of observations discovered, perhaps feeling he has lost out on the chance to "get lucky" and find a high value observation early in the search. (Intuitively, finding a high value observation later is effectively the same as finding a lower value observation earlier because of time discounting.) With these effects added, I anticipate the effect of the population variance on the gap between actual maximum utility and realized utility will be less—the magnitude of the change that occurs will be related to the proportionality constant between a_2 and σ_t.

Varying the rate at which the agent discovers observations may also affect the results found. Depending on the application, it may not be appropriate for the agent to discover observations at a constant rate.[3] True, the constant rate approach works for situations like monogamous dating (where generally one can only 'discover' one person at a time), but imagine a sought-after house hits the market after several months of anticipation. The sellers, who are 'searching' for the highest utility offer, will likely receive an influx of offers initially; more offers will come in as time progresses but at a decreasing rate. Eventually the number of offers should asymptote to 0 or 1 as $t \rightarrow \infty$, with some associated probability of either outcome as each additional time period is added.

Operationally, one can write such a function—perhaps a shifted, discretized decaying exponential, call it $g(t)$—and add $g(t_0)$ elements to the consideration space, ω, at time t_0 to

[3] As an aside, just because I defined it so that the agent discovers one observation per time period, that doesn't mean you couldn't have chosen a different time scale to establish the same parity.

accomplish this. Variants of this model could include either searches where the agent is allowed to return to previously investigated observations or searching where the agent must choose between the offers presented at time t_0. As it is, the constant rate discovery model investigated in this paper is probably not a bad approximation for the behavior that occurs toward the middle or the end of another type of search, like the exponential decay. Trying other functional forms of $g(t)$ may also be fruitful, but such an approach is unlikely to change qualitative behavior. Ideas for this include a simulation of the act of physical searching, i.e. the agent searches outward investigating observations at the circumference of an expanding circle. In this case, because it's the area of the circle that's growing, the number of new observations discovered per circular expansion step should scale as $r(t)^2$, where $r(t)$ is the radius of the circle as a function of time.

Given a more rapidly expanding consideration space, it becomes unreasonable at some point to continually disregard the time it takes the agent to compute $E[x_i]$. Correcting the utility function to account for this might involve something of the form

(15) $$U(\omega,t) = Max(E[\omega]) - a_1(t + T_{rank}) - \frac{a_2}{t},$$

where T_{rank}~$N(\mu_{rank},\sigma^2_{rank})$ and μ_{rank} and σ^2_{rank} are the mean and variance time it takes for the agent to compute $Max(E[\omega])$. This increases the complexity of the model by a significant amount, however, because the agent will also factor this time into his computation of next time period's utility. To do this, he must make assumptions about the distribution of T_{rank} and then revise his estimation of the distribution parameters in a manner similar to the way he revises his estimation of the distribution of Ω. The net effect of these simultaneous processes will likely be heightened sensitivity to initial conditions and thus a more chaotic model, though the exercise is worth the time.

The inclusion of an agent's taste for risk could also add depth to the model. A simple way to accomplish this would be to change the way the agent calculates utility at time $t + 1$ by

(16) $$U_{t+1}[\omega_t, \mu_t, \sigma_t, t] = Max(E[\omega_t + \{\mu_t\}]) \pm a_3(\sigma_t) - a_1(t+1) - \frac{a_2}{t+1},$$

where a_3 is a positive parameter measuring the strength of the agent's taste for or aversion to risk. Note that a risk-averse agent would subtract $a_3(\sigma_t)$ while a risk-loving agent would add the term. Risk here is quantified by the agent's calculated standard deviation of Ω at time t. Another possible interpretation of this term could be an agent's feelings about the future, i.e. optimism vs. pessimism.

Each of these changes, when implemented, has potential to make the current searching model more realistic. At the same time, it is important to implement each individually first to the most basic case before combining them, in order to identify potential difficulties. As complexity in the model increases, I predict increasing sensitivity to parameters and randomly generated initial conditions; this will make interpretation difficult.

But in truth, I don't think I have addressed the biggest hurdle to the model's applicability—of significant concern is its dependence on cardinal ranking. By necessity, in order to compute utility at time t, compute utility at time $t + 1$, or revise the agent's perception of the underlying distribution, each observation must yield a concrete value for $E[x_i]$. Worthy of exploration is the model's generalization to ordinal ranking. While I predict qualitative behavior to be similar, it is presently unclear how one would perform manipulations similar to those in this model to an ordered list of options. This is an avenue worth exploring.

Without a doubt, the questions raised by the results of this model are more numerous and deep than any quantitative relationship developed herein. The reader should instead place emphasis on the qualitative findings—that under the regime founded by current functional forms a rational agent will end a search before utility is truly at a maximum. Increasing the variance in Ω causes the gap between realized utility and true maximum utility to grow. These are the key takeaways.

VI. References:

Augustin, Thomas. "Expected Utility within a Generalized Concept of Probability--a Comprehensive Framework for Decision Making under Ambiguity." *Source Statistical Papers, vol. 43, no. 1, January 2002, pp. 5-22*.

Avrahami, Judith; Guth Werner and Kareev Yaakov. "Games of Competition in a Stochastic Environment." *Source Theory and Decision, vol. 59, no. 4, December 2005, pp. 255-94*.

Blavatskyy, Pavlo R. "Preference Reversals and Probabilistic Decisions." *Source Journal of Risk and Uncertainty, vol. 39, no. 3, December 2009, pp. 237-50*.

Blavatskyy, Pavlo R. "Stochastic Expected Utility Theory." *Source Journal of Risk and Uncertainty, vol. 34, no. 3, June 2007, pp. 259-86*.

Cahuc, Pierre and Andre Zylberberg. 2004. *Labor Economics*. Cambridge: The MIT Press.

Cai, Xiaoqiang Wu and Xianyi Zhou Xian. "Stochastic Scheduling Subject to Preemptive-Repeat Breakdowns with Incomplete Information." *Operations Research, vol. 57, no. 5, September-October 2009, pp. 1236-49*.

Camerer, Colin; Loewenstein George and Weber Martin. "The Curse of Knowledge in Economic Settings: An Experimental Analysis."*: Source Journal of Political Economy, vol. 97, no. 5, October 1989, pp. 1232-54*.

Caputo, Michael R. "The Intrinsic Qualitative Properties of the Classical Optimal Stopping Problem Are Invariant to the Functional Form of the Discount Function." *Natural Resource Modeling, vol. 21, no. 4, Winter 2008, pp. 607-24*.

Cole, Harold and Richard Rogerson. 1999. "Can the Mortensen-Pissarides Matching Model

Match the Business-Cycle Facts?" *Interational Economic Review*, 40(4), pp. 933-59.

Das, Sanmay and John N. Tsitsiklis. "When Is It Important to Know You've Been Rejected? A Search Problem with Probabilistic Appearance of Offers." *Journal of Economic Behavior and Organization, vol. 74, no. 1-2, May 2010, pp. 104-22*.

Diamond, Peter. 1982. "Wage Determination and Efficiency in Search Equilibrium." *The Review of Ecomomic Studies*, 49(2), pp. 217-27.

Diecidue, Enrico; Schmidt Ulrich and Zank Horst. "Parametric Weighting Functions."*: Source Journal of Economic Theory, vol. 144, no. 3, May 2009, pp. 1102-18*.

Evans, Jonathan; Vicky Henderson and David Hobson. "Optimal Timing for an Indivisible Asset Sale." *Mathematical Finance, vol. 18, no. 4, October 2008, pp. 545-67*.

Foster, Dean P. and Vohra Rakesh. "Regret in the on-Line Decision Problem." *Source Games and Economic Behavior, vol. 29, no. 1-2, October-November 1999, pp. 7-35*.

Gapeev, Pavel V. "Discounted Optimal Stopping for Maxima in Diffusion Models with Finite Horizon." *Sonderforschungsbereich 649, Humboldt University, Berlin, Germany, SFB 649 Discussion Papers: SFB649DP2006-057, 2006, 22 pp.*

Grant, Simon; Atsushi Kajii and Ben Polak. "Decomposable Choice under Uncertainty" *Source Journal of Economic Theory, vol. 92, no. 2, June 2000, pp. 169-97*

Golan, Amos; George Judge and Sherman Robinson. "- Recovering Information from Incomplete or Partial Multisectoral Economic Data." - 76(- 3), - 549.

Hey, John D. and Orme Chris. "Investigating Generalizations of Expected Utility Theory Using Experimental Data." *Experiments in environmental economics. Volume 2. 2003, pp. 279-314*.

Horrace, William C. "On Ranking and Selection from Independent Truncated Normal

Distributions."*: Source Journal of Econometrics, vol. 126, no. 2, June 2005, pp. 335-54*.

Howard, David H. "Why Do Transplant Surgeons Turn Down Organs? A Model of the Accept/Reject Decision." *Journal of Health Economics, vol. 21, no. 6, November 2002, pp. 957-69*.

Kondo, Akiko; Kondo Akio; Mannami Yoshihiko and Otsuka Kanako. "Study of Interaction Model by Interaction Purpose on Sightseeing and Business. (in Japanese. With English Summary.)."*: Source Studies in Regional Science, vol. 39, no. 3, 2009, pp. 585-98*.

Kratschmer, Volker and Schoenmakers John. "Representations for Optimal Stopping under Dynamic Monetary Utility Functionals." *Sonderforschungsbereich 649, Humboldt University, Berlin, Germany, SFB 649 Discussion Papers: SFB649DP2009-055, 2009, 26 pp.*

Lempa, Jukka. "On Infinite Horizon Optimal Stopping of General Random Walk." *Mathematical Methods of Operations Research, vol. 67, no. 2, April 2008, pp. 257-68*.

Lipman, Barton L. "Decision Theory without Logical Omniscience: Toward an Axiomatic Framework for Bounded Rationality." *Source Review of Economic Studies, vol. 66, no. 2, April 1999, pp. 339-61*.

Manzini, Paola and Mariotti Marco. "Consumer Choice and Revealed Bounded Rationality." *Economic Theory, vol. 41, no. 3, December 2009, pp. 379-92*.

Miao, Jianjun and Junjian Miao. "Option Exercise with Temptation." *Source Economic Theory, vol. 34, no. 3, March 2008, pp. 473-501*.

Miao, Jianjun and Wang Neng. ":Risk, Uncertainty, and Option Exercise." *Boston University - Department of Economics, Boston University - Department of Economics - Working Papers Series: WP2007-017, 2007, 38 pp.*

Quah, John K. H. and Strulovici Bruno. "Discounting and Patience in Optimal Stopping and

Control Problems." *Northwestern University, Center for Mathematical Studies in Economics and Management Science, Discussion Papers: 1480, 2009*.

Saito, Kota. "A Relationship between Risk and Time Preferences." *Source Northwestern University, Center for Mathematical Studies in Economics and Management Science, Discussion Papers: 1477, 2009*.

Schwartz, Barry. "The Paradox of Choice: Why More Is Less."*: Source New York: HarperCollins, Ecco, 2004, pp. xi, 265*.

Stadje, Wolfgang. "An Optimal Stopping Problem with Two Levels of Incomplete Information." *Mathematical Methods of Operations Research, vol. 45, no. 1, 1997, pp. 119-31*.

VII. Tables:

Table 1: Premature search data. Parameters: $a_1 = 0.05$, $a_2 = 5.0$

σ/μ Ratio	Δt (in periods) from theoretical value	Uncertainty
0.00	0.00	0.00
0.003125	1.56	0.77
0.00625	2.60	0.87
0.0125	3.56	0.51
0.025	4.60	0.76
0.05	5.52	1.00
0.0625	5.64	0.70
0.075	5.80	0.58
0.10	5.88	0.67
0.125	6.56	0.51
0.15	6.64	0.76

Table 2: Data supporting regression of percent decrease in realized utility vs. percent increase in σ/μ.

Percent Increase in σ/μ	Percent Decrease in U
0.00	0.00
1.25	4.98
2.50	10.52
3.75	14.45
5.00	16.22
6.25	22.23
7.50	27.14
8.75	32.60
10.00	36.10
11.25	40.82
12.50	47.68

VIII. Figures:

Figure 1: Maximum utility occurs when time disutility balances alleviation of regret.

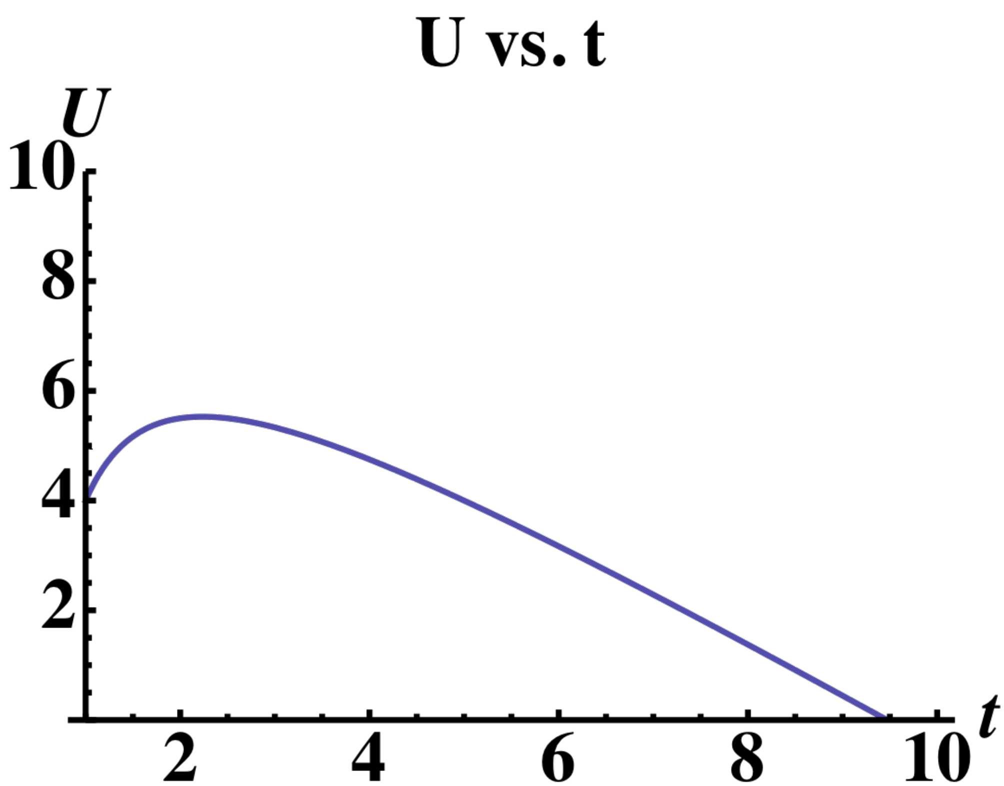

Figure 2: An agent's prior knowledge is randomly generated based on the underlying distribution of Ω.

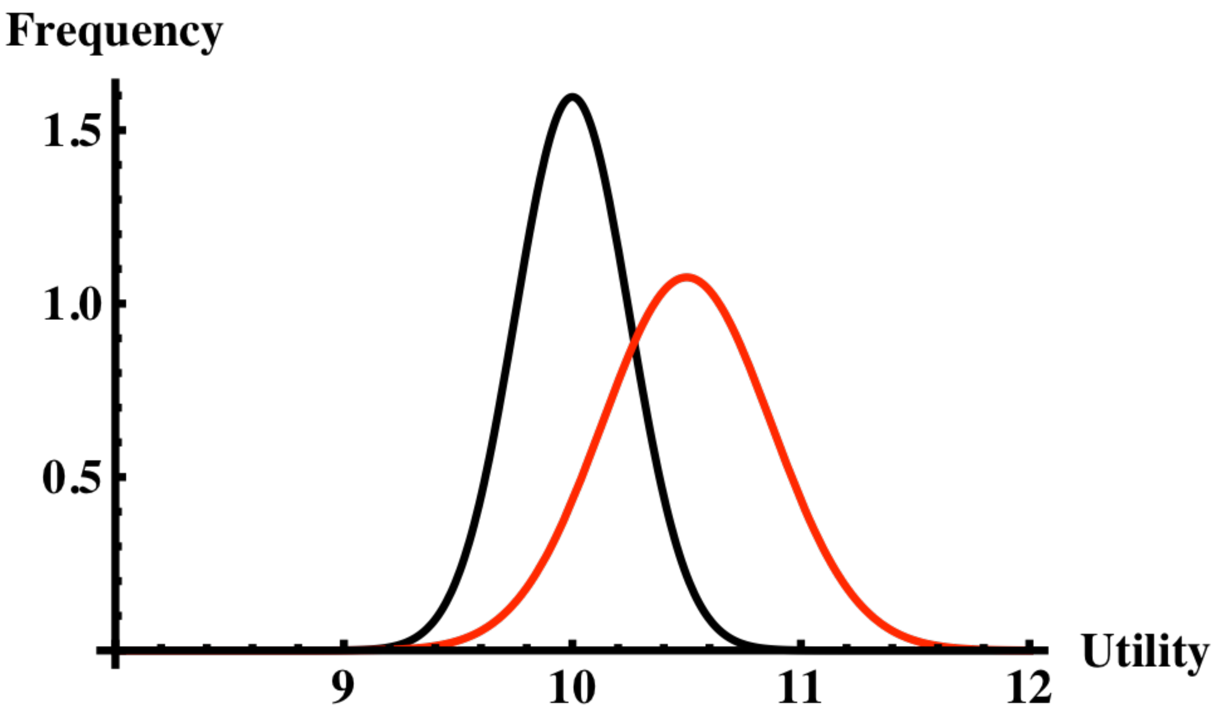

Figure 3: Acquisition of the first observation and subsequent revision.

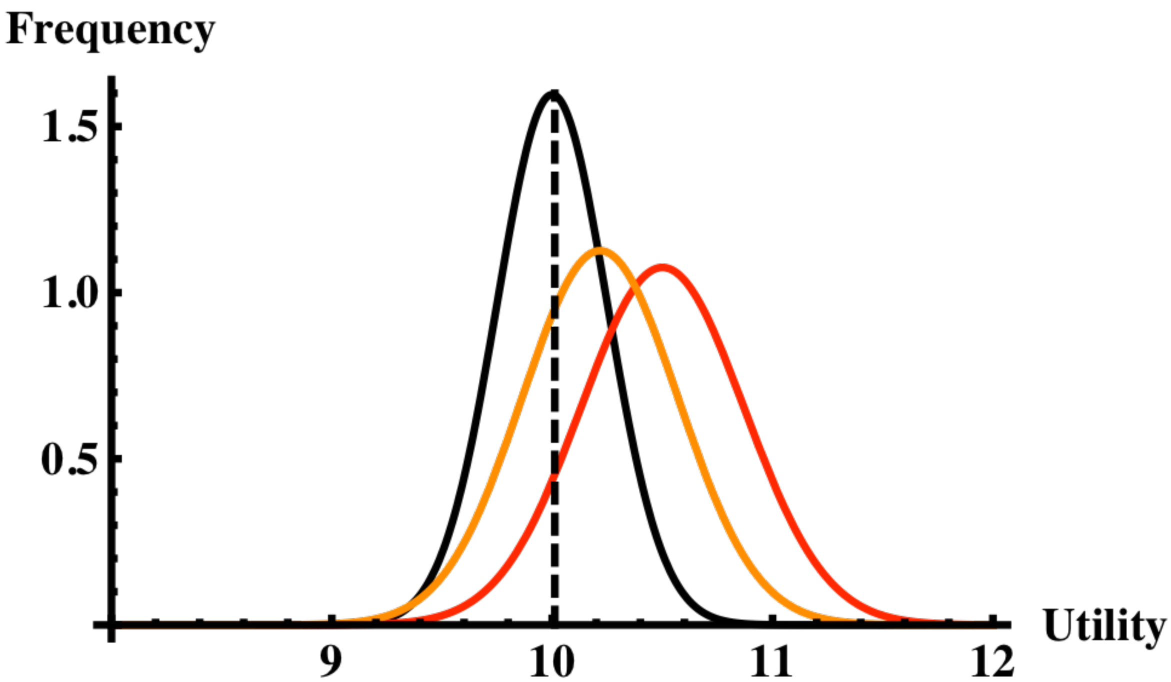

Figure 4: Convergence of agent's prior knowledge to actual distribution of Ω, $N_0 = 10$.

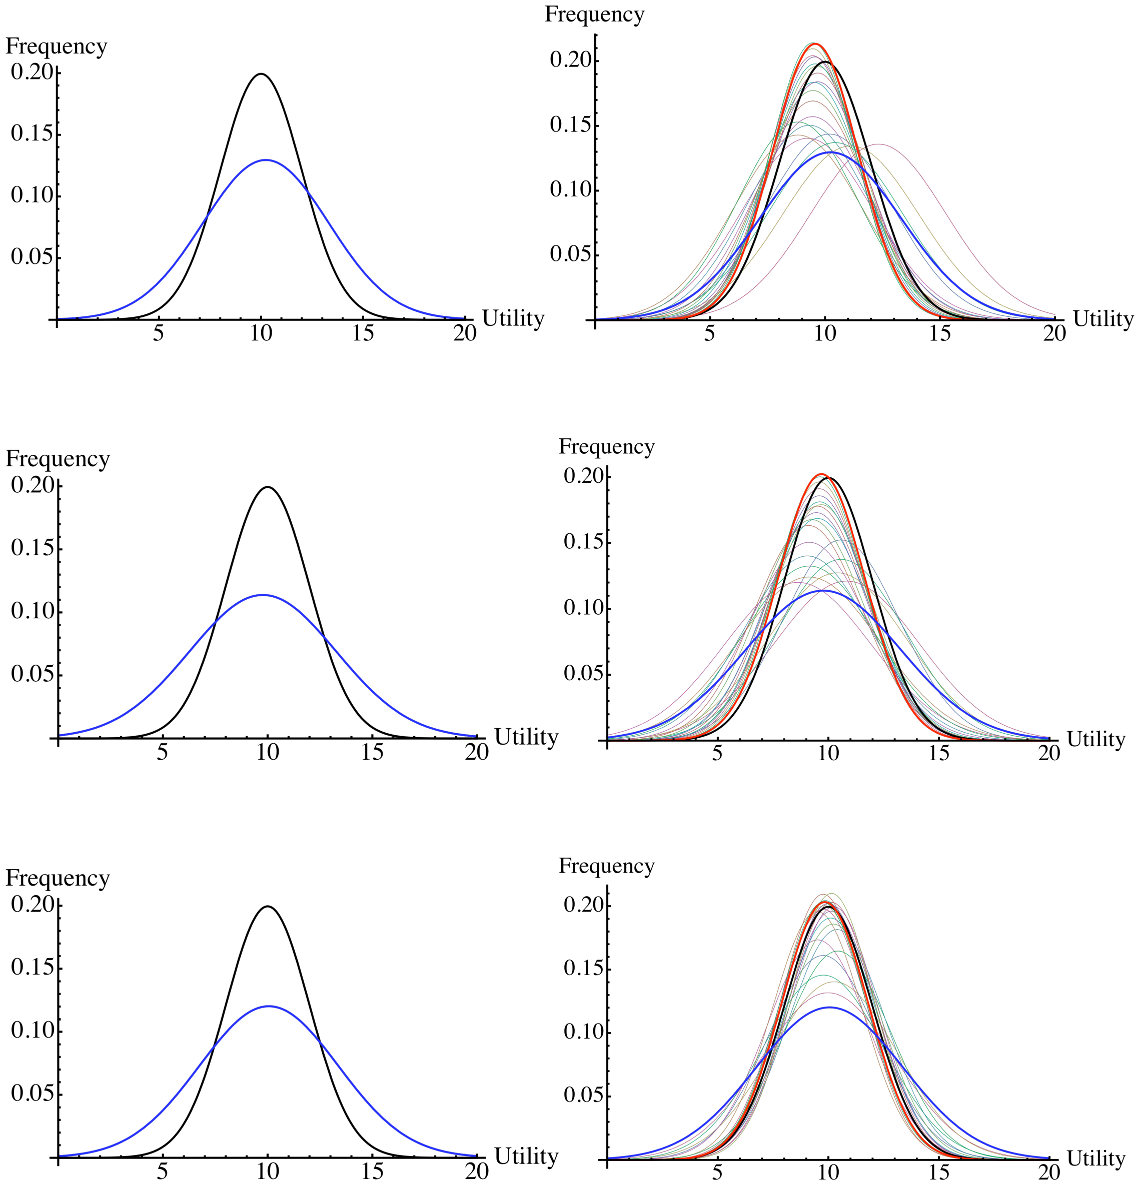

Figure 5: Convergence of agent's prior knowledge to actual distribution of Ω, $N_0 = 5$.

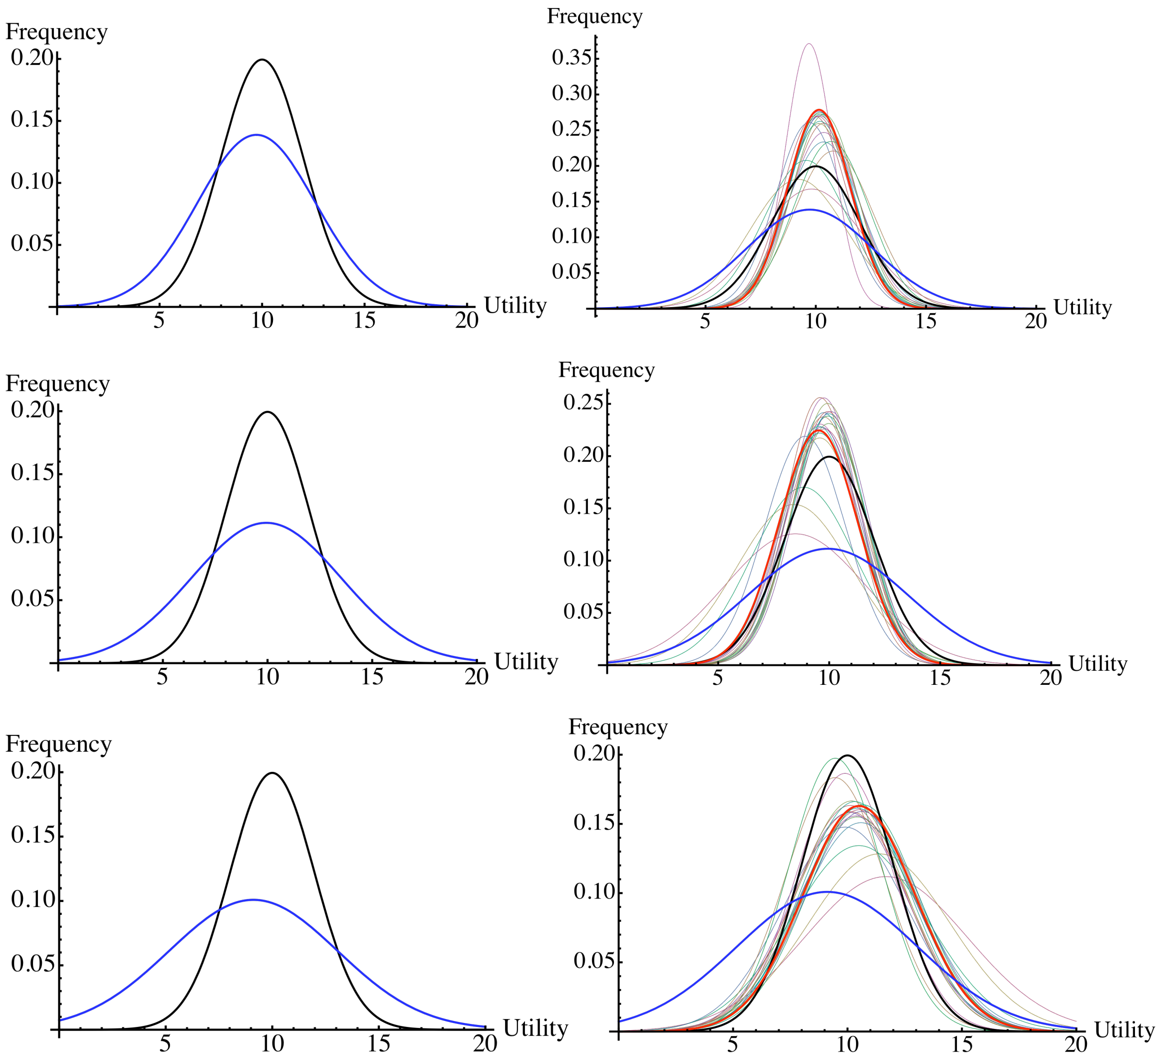

Figure 6: Utility vs. Time. Expected utility of $t + 1$ is graphed in red, utility at time t is graphed in blue.

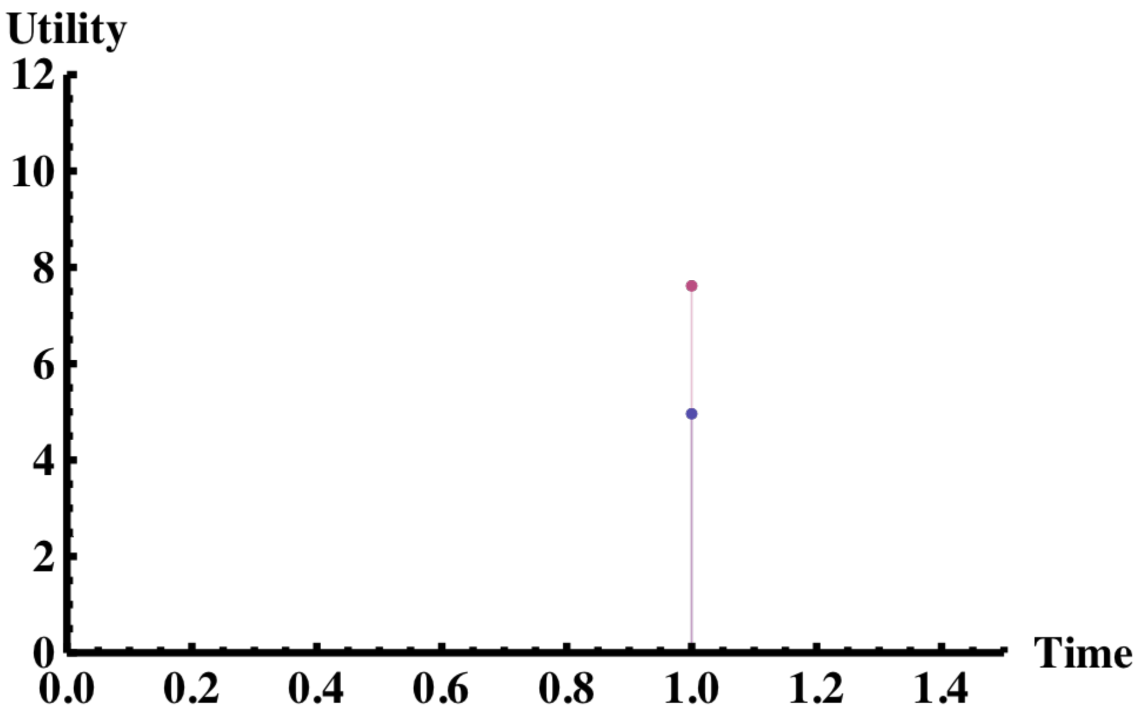

Figure 7: Agent's conception of distribution of Ω graphed in magenta at the conclusion of the search initialized in Figure 3.

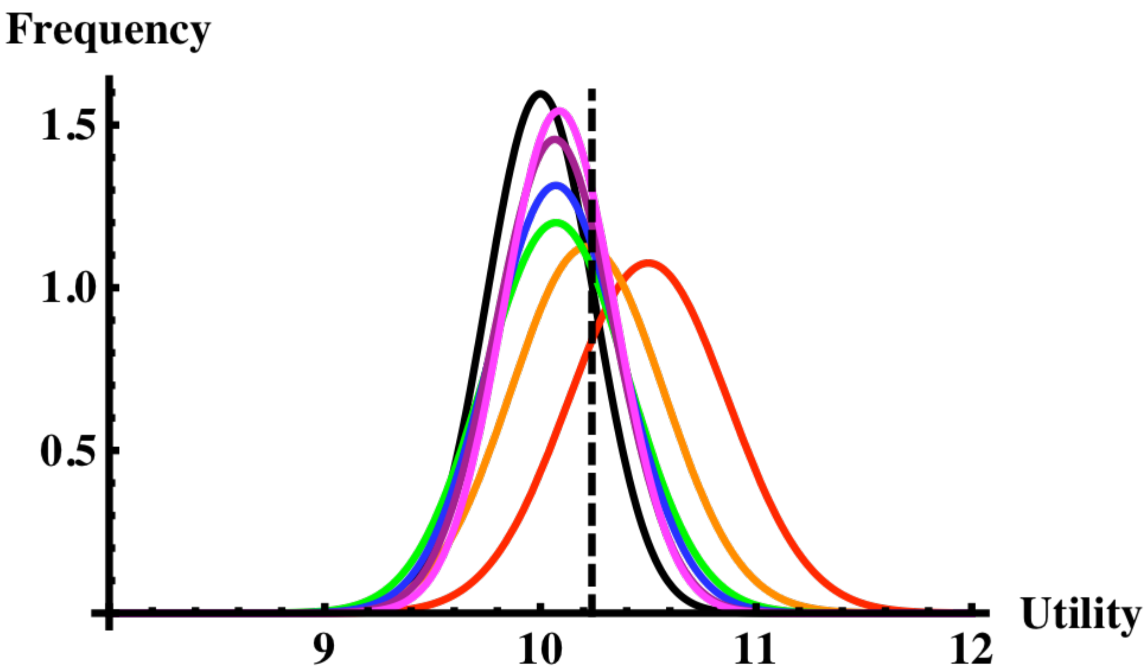

Figure 8: Expected utility and realized utility converge at time $t = 5$ for search initialized in Figure 3

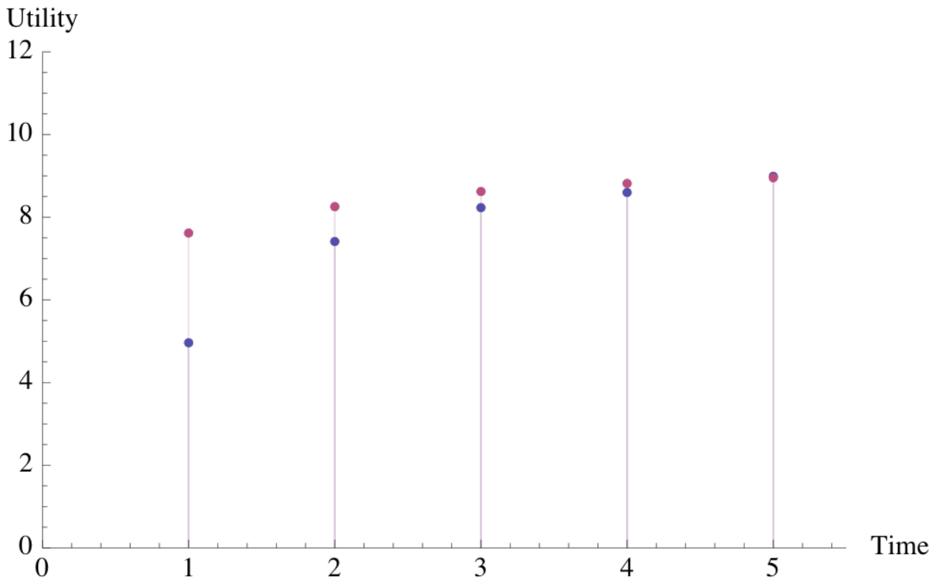

Figure 9: Realized utility could be improved if the agent searches past his rational stopping point

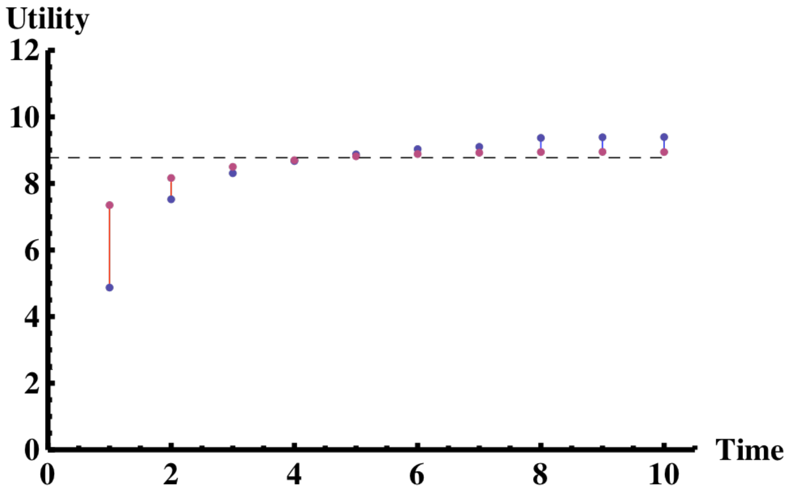

Figure 10: Search with variance stops before the no-variance calculated optimal point

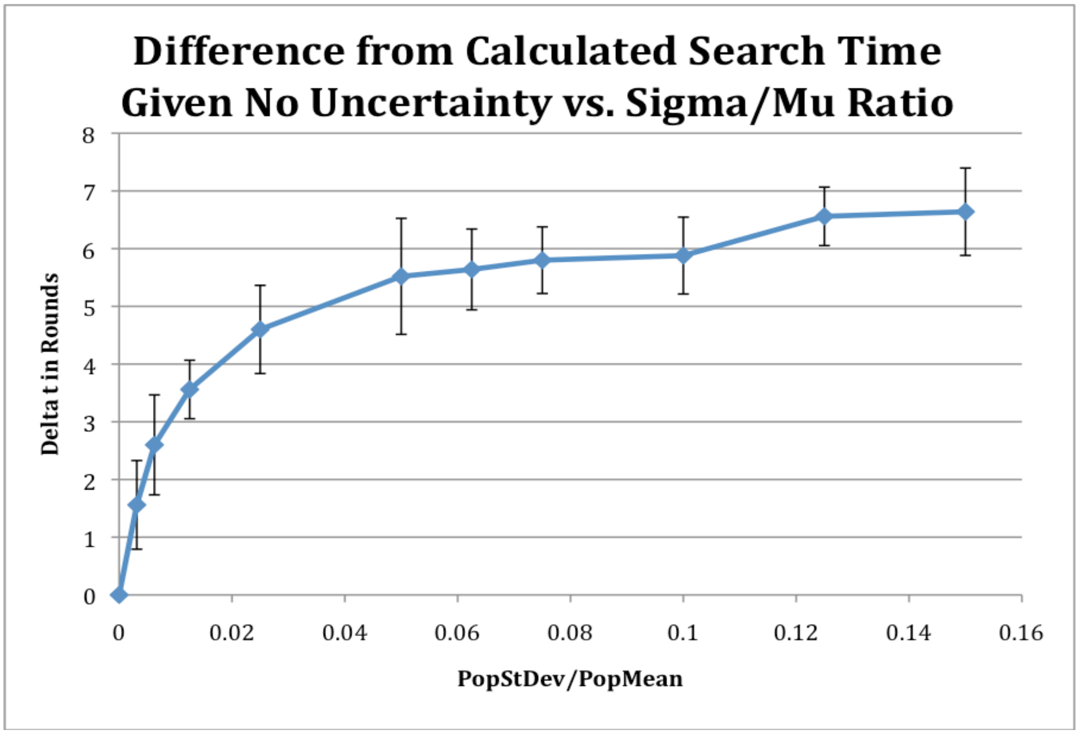

Figure 11: The gap between true maximum utility and stopping utility grows with variance in Ω

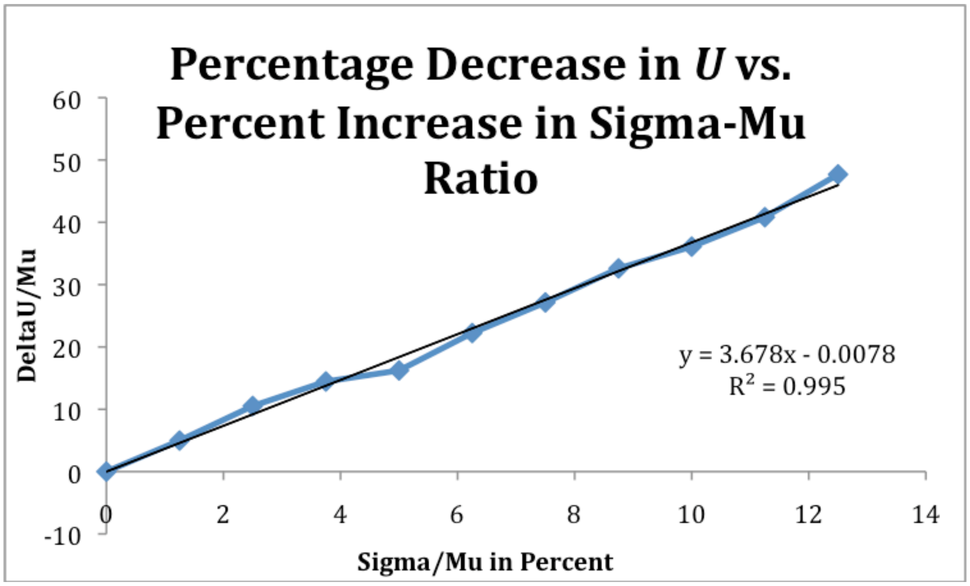

Figure 12: The 'gap' between true maximum utility and realized stopping utility

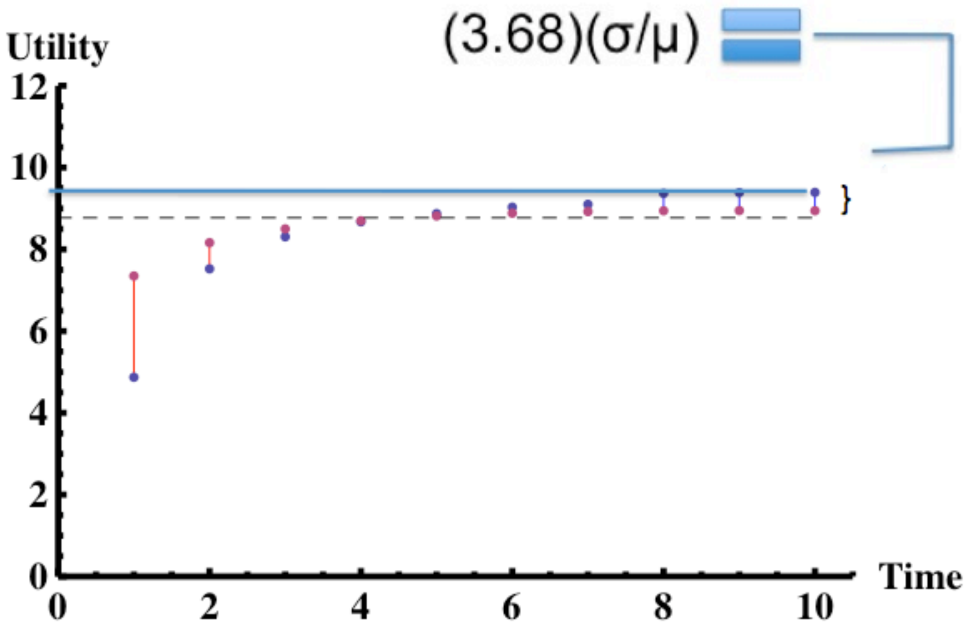

Figure 13: The agent must rationally settle for utility less than that which he once had realized

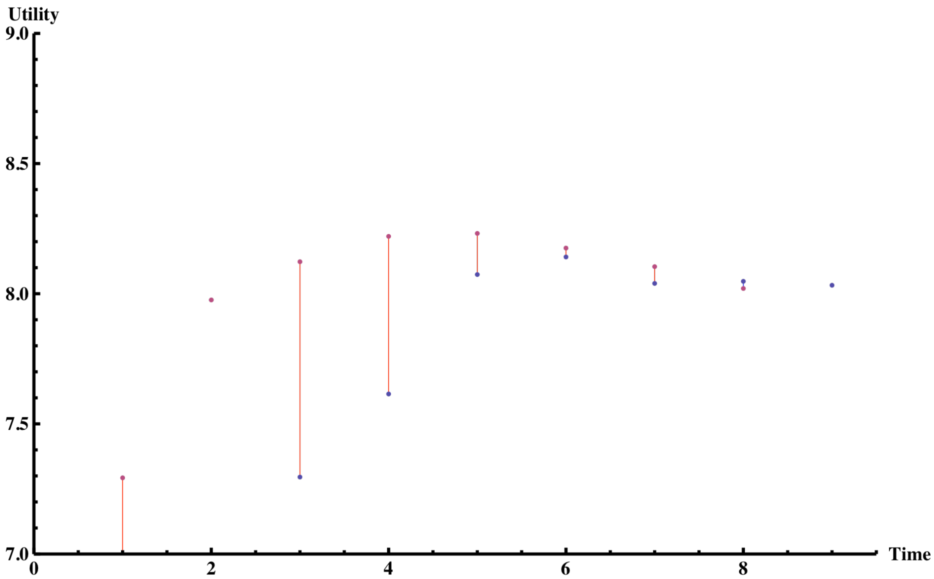

Figure 14: A full simulation run with code.

David Menasche
11/28/10

Clean Run 2. Graphics Included.

μ = 10;
σ = .25;

$$\text{Plot}\left[\frac{1}{\sqrt{2\,\pi\,\sigma^2}}\,e^{\frac{-(x-\mu)^2}{2\,\sigma^2}},\ \{x,\ 8,\ 12\},\ \text{PlotRange}\to\text{All},\ \text{PlotStyle}\to\{\text{Black},\ \text{Thick}\},\ \text{AxesLabel}\to\{\text{"Utility"},\ \text{"Frequency"}\},\ \text{LabelStyle}\to\{\text{Bold},\ \text{Medium}\},\ \text{AxesStyle}\to\text{Thick}\right]$$

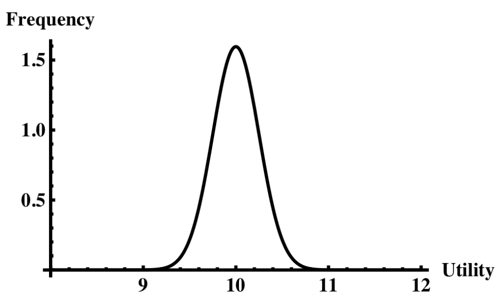

```
ω = {RandomReal[NormalDistribution[μ, σ]]};
n = Length[ω];
No = 10;
μprior = μ * RandomReal[{.9, 1.1}];
σprior = σ * RandomReal[{1, 2}];
```

$$\text{Plot}\left[\left\{\frac{1}{\sqrt{2\,\pi\,\sigma^2}}\,e^{\frac{-(x-\mu)^2}{2\,\sigma^2}},\ \frac{1}{\sqrt{2\,\pi\,\sigma\text{prior}^2}}\,e^{\frac{-(x-\mu\text{prior})^2}{2\,\sigma\text{prior}^2}}\right\},\ \{x,\ 8,\ 12\},\right.$$

```
 PlotRange → All, PlotStyle → {{Black, Thick}, {Red, Thick}},
 AxesLabel → {"Utility", "Frequency"}, LabelStyle → {Bold, Medium}, AxesStyle → Thick]
μpriorlist = {μprior}
σpriorlist = {σprior}
```

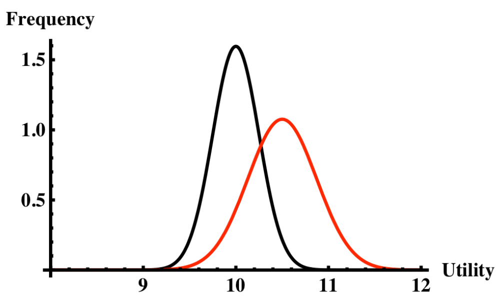


```
{10.5003}

{0.370832}
```

```
a1 = .05;
a2 = 5;
T = Table[i, {i, 1000}];
t = T[[1]];
o = Length[ω];
```

$$\text{Utility}[\omega\_] := \text{Max}[\omega] - \text{a1} * \text{T}[[\text{t}]] - \frac{\text{a2}}{\text{n}};$$

$$\text{NextRoundUtility}[\mu\text{prior\_},\ \text{t\_},\ \text{n\_}] := \mu\text{prior} - \text{a1}\ (\text{T}[[\text{t}+1]]) - \frac{\text{a2}}{\text{n}+1}$$

$$\texttt{sol = Solve}\left[\left\{\mu\texttt{est} == \left((\texttt{n} * \mu\texttt{priorlist[[n]]}) + \left(\texttt{Abs}\left[\frac{\omega\texttt{[[n]]} - \mu\texttt{priorlist[[n]]}}{\sigma\texttt{est}}\right] * \omega\texttt{[[n]]}\right)\right) \Big/ \left(\texttt{n} + \texttt{Abs}\left[\frac{\omega\texttt{[[n]]} - \mu\texttt{priorlist[[n]]}}{\sigma\texttt{est}}\right]\right),\right.\right.$$

$$\left.\left.\sigma\texttt{est} == \frac{\texttt{No - n}}{\texttt{No}}\,\sigma\texttt{prior} + \frac{\texttt{n}}{\texttt{No}}\sqrt{\frac{1}{\texttt{n}}\sum_{\texttt{i}=1}^{\texttt{n}}(\omega\texttt{[[i]]} - \mu\texttt{est})^2}\right\}, \{\mu\texttt{est}, \sigma\texttt{est}\}\right]$$

```
μpriorlist = Join[μpriorlist, {μest /. sol[[1]]}]
σpriorlist = Join[σpriorlist, {σest /. sol[[1]]}]
```

```
{{σest → 0.354283, μest → 10.2172}}

{10.5003, 10.2172}

{0.370832, 0.354283}
```

```
l1 = Line[{{ω[[o]], 0}, {ω[[o]], 1.6}}];
```

$$\texttt{Show}\left[\texttt{Plot}\left[\left\{\frac{1}{\sqrt{2\pi\sigma^2}}\,e^{\frac{-(x-\mu)^2}{2\sigma^2}}, \frac{1}{\sqrt{2\pi\,\sigma\texttt{priorlist[[1]]}^2}}\,e^{\frac{-(x-\mu\texttt{priorlist[[1]]})^2}{2\,\sigma\texttt{priorlist[[1]]}^2}}, \frac{1}{\sqrt{2\pi\,\sigma\texttt{priorlist[[2]]}^2}}\,e^{\frac{-(x-\mu\texttt{priorlist[[2]]})^2}{2\,\sigma\texttt{priorlist[[2]]}^2}}\right\},\right.\right.$$

```
  {x, 8, 12}, PlotRange → All, PlotStyle → {{Black, Thick}, {Red, Thick}, {Orange, Thick}}],
 Graphics[{Thick, Dashed, l1}], AxesLabel → {"Utility", "Frequency"},
 LabelStyle → {Bold, Medium}, AxesStyle → Thick]
```

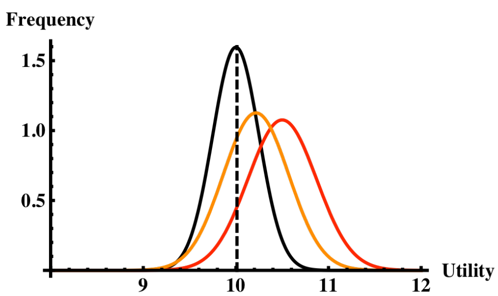


```
U1 = Utility[ω]
U1next = NextRoundUtility[μpriorlist[[Length[μpriorlist]]], t, n]
```

```
4.96183

7.61718
```

```
Ulist = {{t, U1}};
Unextlist = {{t, U1next}};
```

```
ListPlot[{Ulist, Unextlist}, PlotRange → {{0, 1.5}, {0, 12}}, AxesLabel → {"Time", "Utility"},
 Filling → Axis, LabelStyle → {Bold, Medium}, AxesStyle → Thick]
```

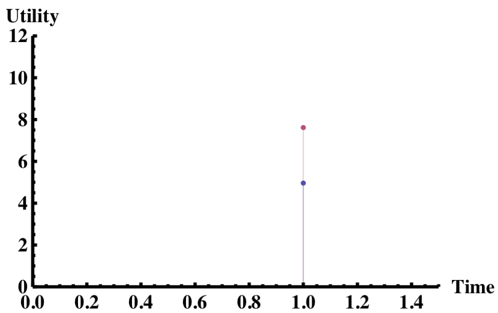


```
ω = Join[ω, {RandomReal[NormalDistribution[μ, σ]]}]
n = 2;
t = t + 1;
o = o + 1;
```

{10.0118, 9.82632}

$$\texttt{sol = Solve}\Big[\Big\{\mu\texttt{est} == \left((\texttt{n} * \mu\texttt{priorlist[[n]]}) + \left(\texttt{Abs}\left[\frac{\omega\texttt{[[n]]} - \mu\texttt{priorlist[[n]]}}{\sigma\texttt{est}}\right] * \omega\texttt{[[n]]}\right)\right) \Big/ \left(\texttt{n} + \texttt{Abs}\left[\frac{\omega\texttt{[[n]]} - \mu\texttt{priorlist[[n]]}}{\sigma\texttt{est}}\right]\right),$$

$$\sigma\texttt{est} == \frac{\texttt{No - n}}{\texttt{No}}\sigma\texttt{prior} + \frac{\texttt{n}}{\texttt{No}}\sqrt{\frac{1}{\texttt{n}}\sum_{\texttt{i}=1}^{\texttt{n}}(\omega\texttt{[[i]]} - \mu\texttt{est})^2}\Big\}, \{\mu\texttt{est}, \sigma\texttt{est}\}\Big]$$

```
μpriorlist = Join[μpriorlist, {μest /. sol[[1]]}]
σpriorlist = Join[σpriorlist, {σest /. sol[[1]]}]
```

{{σest → 0.332522, μest → 10.0725}}

{10.5003, 10.2172, 10.0725}

{0.370832, 0.354283, 0.332522}

```
l2 = Line[{{ω[[o]], 0}, {ω[[o]], 1.6}}];
Show[Plot[{1/Sqrt[2 π σ^2] e^(-(x-μ)^2/(2 σ^2)), 1/Sqrt[2 π σpriorlist[[1]]^2] e^(-(x-μpriorlist[[1]])^2/(2 σpriorlist[[1]]^2)),
   1/Sqrt[2 π σpriorlist[[2]]^2] e^(-(x-μpriorlist[[2]])^2/(2 σpriorlist[[2]]^2)), 1/Sqrt[2 π σpriorlist[[3]]^2] e^(-(x-μpriorlist[[3]])^2/(2 σpriorlist[[3]]^2))}, {x, 8, 12},
  PlotRange → All, PlotStyle → {{Black, Thick}, {Red, Thick}, {Orange, Thick}, {Green, Thick}}],
 Graphics[{Dashed, Thick, l2}], AxesLabel → {"Utility", "Frequency"},
 LabelStyle → {Bold, Medium}, AxesStyle → Thick]
```

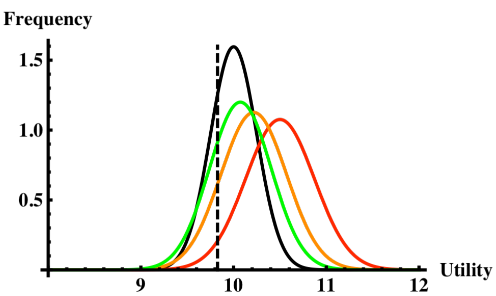


```
U2 = Utility[ω]
U2next = NextRoundUtility[μpriorlist[[Length[μpriorlist]]], t, n]
```

```
7.41183
```

```
8.25583
```

```
Ulist = Join[Ulist, {{t, U2}}];
Unextlist = Join[Unextlist, {{t, U2next}}];
```

```
ListPlot[{Ulist, Unextlist}, PlotRange → {{0, 2.5}, {0, 12}}, AxesLabel → {"Time", "Utility"},
 Filling → Axis, LabelStyle → {Bold, Medium}, AxesStyle → Thick]
```

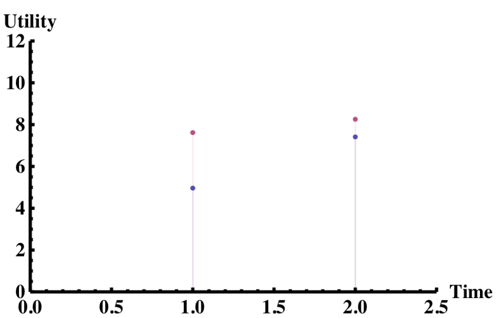

```
ω = Join[ω, {RandomReal[NormalDistribution[μ, σ]]}]
n = 3;
t = t + 1;
o = o + 1;
```

{10.0118, 9.82632, 10.0479}

$$\texttt{sol = Solve}\left[\left\{\mu\texttt{est} == \left((\texttt{n} * \mu\texttt{priorlist[[n]]}) + \left(\texttt{Abs}\left[\frac{\omega\texttt{[[n]]} - \mu\texttt{priorlist[[n]]}}{\sigma\texttt{est}}\right] * \omega\texttt{[[n]]}\right)\right) \Big/ \left(\texttt{n} + \texttt{Abs}\left[\frac{\omega\texttt{[[n]]} - \mu\texttt{priorlist[[n]]}}{\sigma\texttt{est}}\right]\right),\right.\right.$$

$$\left.\left.\sigma\texttt{est} == \frac{\texttt{No - n}}{\texttt{No}}\,\sigma\texttt{prior} + \frac{\texttt{n}}{\texttt{No}}\sqrt{\frac{1}{\texttt{n}}\sum_{\texttt{i=1}}^{\texttt{n}}(\omega\texttt{[[i]]} - \mu\texttt{est})^2}\right\}, \{\mu\texttt{est}, \sigma\texttt{est}\}\right]$$

```
μpriorlist = Join[μpriorlist, {μest /. sol[[1]]}]
σpriorlist = Join[σpriorlist, {σest /. sol[[1]]}]
```

{{σest → 0.303558, μest → 10.0718}}

{10.5003, 10.2172, 10.0725, 10.0718}

{0.370832, 0.354283, 0.332522, 0.303558}

```
l3 = Line[{{ω[[o]], 0}, {ω[[o]], 1.5}}];
```

$$\texttt{Show}\left[\texttt{Plot}\left[\left\{\frac{1}{\sqrt{2\pi\sigma^2}}e^{\frac{-(x-\mu)^2}{2\sigma^2}}, \frac{1}{\sqrt{2\pi\,\sigma\texttt{priorlist[[1]]}^2}}e^{\frac{-(x-\mu\texttt{priorlist[[1]]})^2}{2\,\sigma\texttt{priorlist[[1]]}^2}},\right.\right.\right.$$

$$\frac{1}{\sqrt{2\pi\,\sigma\texttt{priorlist[[2]]}^2}}e^{\frac{-(x-\mu\texttt{priorlist[[2]]})^2}{2\,\sigma\texttt{priorlist[[2]]}^2}}, \frac{1}{\sqrt{2\pi\,\sigma\texttt{priorlist[[3]]}^2}}e^{\frac{-(x-\mu\texttt{priorlist[[3]]})^2}{2\,\sigma\texttt{priorlist[[3]]}^2}},$$

$$\left.\frac{1}{\sqrt{2\pi\,\sigma\texttt{priorlist[[4]]}^2}}e^{\frac{-(x-\mu\texttt{priorlist[[4]]})^2}{2\,\sigma\texttt{priorlist[[4]]}^2}}\right\}, \{x, 8, 12\}, \texttt{PlotRange} \to \texttt{All},$$

```
  PlotStyle → {{Black, Thick}, {Red, Thick}, {Orange, Thick}, {Green, Thick}, {Blue, Thick}}],
 Graphics[{Dashed, Thick, l3}], AxesLabel → {"Utility", "Frequency"},
 LabelStyle → {Bold, Medium}, AxesStyle → Thick]
```

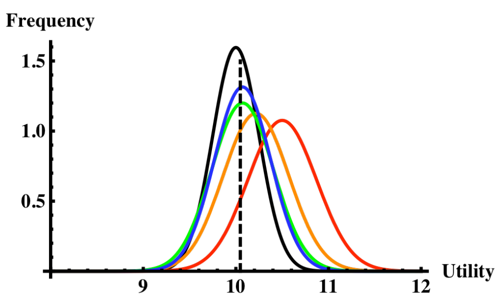

```
U3 = Utility[ω]
U3next = NextRoundUtility[μpriorlist[[Length[μpriorlist]]], t, n]
```

```
8.23121
```

```
8.62185
```

```
Ulist = Join[Ulist, {{t, U3}}];
Unextlist = Join[Unextlist, {{t, U3next}}];
```

```
ListPlot[{Ulist, Unextlist}, PlotRange → {{0, 3.5}, {0, 12}},
 AxesLabel → {"Time", "Utility"}, Filling → Axis]
```

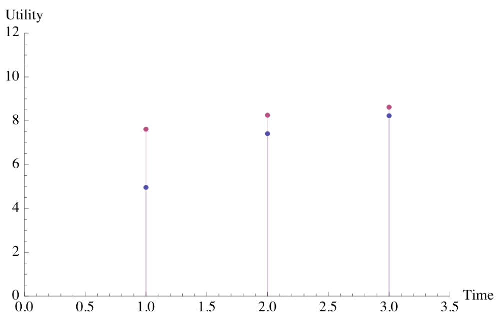


```
ω = Join[ω, {RandomReal[NormalDistribution[μ, σ]]}]
n = 4;
t = t + 1;
o = o + 1;
```

```
{10.0118, 9.82632, 10.0479, 9.99313}
```

$$\text{sol} = \text{Solve}\Big[\Big\{\mu\text{est} == \left((\text{n} * \mu\text{priorlist}[[\text{n}]]) + \left(\text{Abs}\left[\frac{\omega[[\text{n}]] - \mu\text{priorlist}[[\text{n}]]}{\sigma\text{est}}\right] * \omega[[\text{n}]]\right)\right) \Big/ \left(\text{n} + \text{Abs}\left[\frac{\omega[[\text{n}]] - \mu\text{priorlist}[[\text{n}]]}{\sigma\text{est}}\right]\right),$$

$$\sigma\text{est} == \frac{\text{No} - \text{n}}{\text{No}}\,\sigma\text{prior} + \frac{\text{n}}{\text{No}}\sqrt{\frac{1}{\text{n}}\sum_{\text{i}=1}^{\text{n}}(\omega[[\text{i}]] - \mu\text{est})^2}\Big\}, \{\mu\text{est}, \sigma\text{est}\}\Big]$$

```
μpriorlist = Join[μpriorlist, {μest /. sol[[1]]}]
σpriorlist = Join[σpriorlist, {σest /. sol[[1]]}]
```

```
{{σest → 0.27406, μest → 10.0666}}
```

```
{10.5003, 10.2172, 10.0725, 10.0718, 10.0666}
```

```
{0.370832, 0.354283, 0.332522, 0.303558, 0.27406}
```

```
l4 = Line[{{ω[[o]], 0}, {ω[[o]], 1.5}}];
Show[Plot[{1/Sqrt[2 π σ^2] e^(-(x-μ)^2/(2 σ^2)), 1/Sqrt[2 π σpriorlist[[1]]^2] e^(-(x-μpriorlist[[1]])^2/(2 σpriorlist[[1]]^2)),
    1/Sqrt[2 π σpriorlist[[2]]^2] e^(-(x-μpriorlist[[2]])^2/(2 σpriorlist[[2]]^2)), 1/Sqrt[2 π σpriorlist[[3]]^2] e^(-(x-μpriorlist[[3]])^2/(2 σpriorlist[[3]]^2)),
    1/Sqrt[2 π σpriorlist[[4]]^2] e^(-(x-μpriorlist[[4]])^2/(2 σpriorlist[[4]]^2)), 1/Sqrt[2 π σpriorlist[[5]]^2] e^(-(x-μpriorlist[[5]])^2/(2 σpriorlist[[5]]^2))},
  {x, 8, 12}, PlotRange → All, PlotStyle → {{Black, Thick}, {Red, Thick}, {Orange, Thick},
    {Green, Thick}, {Blue, Thick}, {Purple, Thick}}, AxesLabel → {"Utility", "Frequency"}],
 Graphics[{Thick, Dashed, l4}], LabelStyle → {Bold, Medium}, AxesStyle → Thick]
```

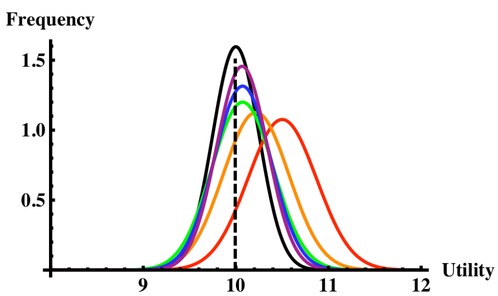


```
U4 = Utility[ω]
U4next = NextRoundUtility[μpriorlist[[Length[μpriorlist]]], t, n]
```

8.59788

8.81657

```
Ulist = Join[Ulist, {{t, U4}}];
Unextlist = Join[Unextlist, {{t, U4next}}];
```

```
ListPlot[{Ulist, Unextlist}, PlotRange → {{0, 4.5}, {0, 12}}, AxesLabel → {"Time", "Utility"},
 Filling → Axis, LabelStyle → {Bold, Medium}, AxesStyle → Thick]
```

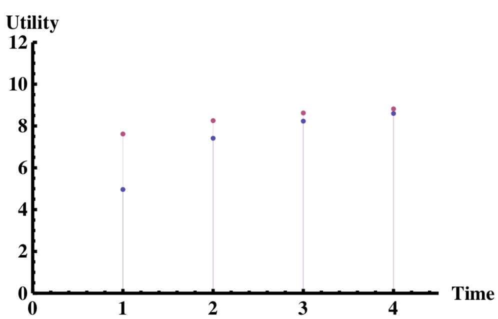

```
ω = Join[ω, {RandomReal[NormalDistribution[μ, σ]]}]
n = 5;
t = t + 1;
o = o + 1;
```

{10.0118, 9.82632, 10.0479, 9.99313, 10.2387}

$$\text{sol} = \text{Solve}\left[\left\{\mu\text{est} == \left((n * \mu\text{priorlist}[[n]]) + \left(\text{Abs}\left[\frac{\omega[[n]] - \mu\text{priorlist}[[n]]}{\sigma\text{est}}\right] * \omega[[n]]\right)\right) \Big/ \left(n + \text{Abs}\left[\frac{\omega[[n]] - \mu\text{priorlist}[[n]]}{\sigma\text{est}}\right]\right), \sigma\text{est} == \frac{\text{No} - n}{\text{No}}\,\sigma\text{prior} + \frac{n}{\text{No}}\sqrt{\frac{1}{n}\sum_{i=1}^{n}(\omega[[i]] - \mu\text{est})^2}\right\}, \{\mu\text{est}, \sigma\text{est}\}\right]$$

```
μpriorlist = Join[μpriorlist, {μest /. sol[[1]]}]
σpriorlist = Join[σpriorlist, {σest /. sol[[1]]}]
```

{{σest → 0.2585, μest → 10.0868}}

{10.5003, 10.2172, 10.0725, 10.0718, 10.0666, 10.0868}

{0.370832, 0.354283, 0.332522, 0.303558, 0.27406, 0.2585}

```
l5 = Line[{{ω[[o]], 0}, {ω[[o]], 1.6}}];
```

$$\text{Show}\left[\text{Plot}\left[\left\{\frac{1}{\sqrt{2\pi\sigma^2}}e^{\frac{-(x-\mu)^2}{2\sigma^2}}, \frac{1}{\sqrt{2\pi\,\sigma\text{priorlist}[[1]]^2}}e^{\frac{-(x-\mu\text{priorlist}[[1]])^2}{2\,\sigma\text{priorlist}[[1]]^2}}, \frac{1}{\sqrt{2\pi\,\sigma\text{priorlist}[[2]]^2}}e^{\frac{-(x-\mu\text{priorlist}[[2]])^2}{2\,\sigma\text{priorlist}[[2]]^2}}, \frac{1}{\sqrt{2\pi\,\sigma\text{priorlist}[[3]]^2}}e^{\frac{-(x-\mu\text{priorlist}[[3]])^2}{2\,\sigma\text{priorlist}[[3]]^2}}, \frac{1}{\sqrt{2\pi\,\sigma\text{priorlist}[[4]]^2}}e^{\frac{-(x-\mu\text{priorlist}[[4]])^2}{2\,\sigma\text{priorlist}[[4]]^2}}, \frac{1}{\sqrt{2\pi\,\sigma\text{priorlist}[[5]]^2}}e^{\frac{-(x-\mu\text{priorlist}[[5]])^2}{2\,\sigma\text{priorlist}[[5]]^2}}, \frac{1}{\sqrt{2\pi\,\sigma\text{priorlist}[[6]]^2}}e^{\frac{-(x-\mu\text{priorlist}[[6]])^2}{2\,\sigma\text{priorlist}[[6]]^2}}\right\}, \{x, 8, 12\}, \text{PlotRange} \to \text{All},\right.$$

```
  PlotStyle → {{Black, Thick}, {Red, Thick}, {Orange, Thick}, {Green, Thick},
    {Blue, Thick}, {Purple, Thick}, {Magenta, Thick}}, LabelStyle → {Bold, Medium},
  AxesStyle → Thick, AxesLabel → {"Utility", "Frequency"}], Graphics[{Thick, Dashed, l5}]]
```

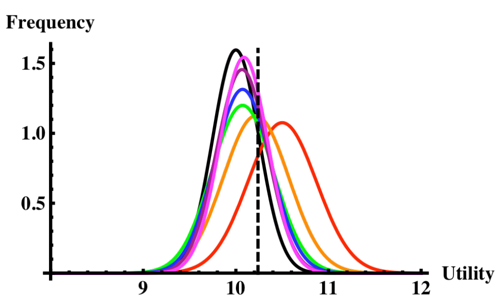

```
U5 = Utility[ω]
U5next = NextRoundUtility[μpriorlist[[Length[μpriorlist]]], t, n]
```

8.98868

8.95347

```
Ulist = Join[Ulist, {{t, U5}}];
Unextlist = Join[Unextlist, {{t, U5next}}];
```

```
ListPlot[{Ulist, Unextlist}, PlotRange → {{0, 5.5}, {0, 12}},
 AxesLabel → {"Time", "Utility"}, Filling → Axis]
```

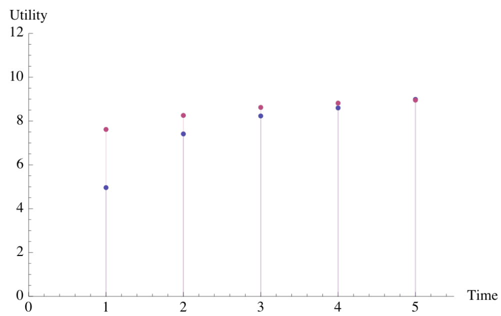

The search would end here.

```
ω = Join[ω, {RandomReal[NormalDistribution[μ, σ]]}]
n = 6;
t = t + 1;
o = o + 1;
```

{10.0118, 9.82632, 10.0479, 9.99313, 10.2387, 9.95512}

```
U6 = Utility[ω]
U6next = NextRoundUtility[μpriorlist[[Length[μpriorlist]]], t, n]
```

9.10535

9.02251

```
Ulist = Join[Ulist, {{t, U6}}];
Unextlist = Join[Unextlist, {{t, U6next}}];
```

```
ListPlot[{Ulist, Unextlist}, PlotRange → {{0, 6.5}, {0, 12}}, AxesLabel → {"Time", "Utility"},
 Filling → Axis, LabelStyle → {Bold, Medium}, AxesStyle → Thick]
```

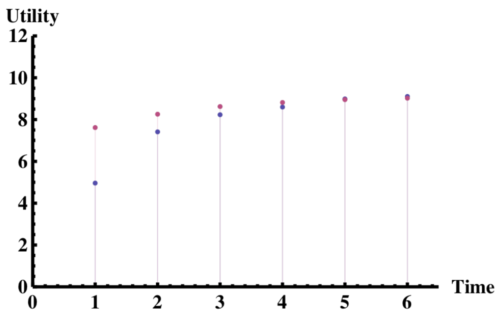


```
ω = Join[ω, {RandomReal[NormalDistribution[μ, σ]]}]
n = 7;
t = t + 1;
o = o + 1;
```

{10.0118, 9.82632, 10.0479, 9.99313, 10.2387, 9.95512, 10.2784}

```
U7 = Utility[ω]
U7next = NextRoundUtility[μpriorlist[[Length[μpriorlist]]], t, n]
```

9.21415

9.0618

```
Ulist = Join[Ulist, {{t, U7}}];
Unextlist = Join[Unextlist, {{t, U7next}}];
```

```
ListPlot[{Ulist, Unextlist}, PlotRange → {{0, 7.5}, {0, 12}}, AxesLabel → {"Time", "Utility"},
 AxesStyle → Thick, LabelStyle → {Bold, Medium}, Filling → {1 → {{2}, {Red, Blue}}}]
```

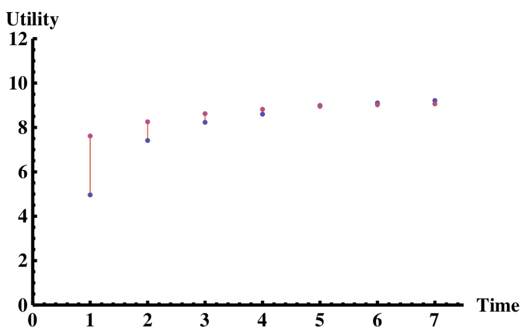


```
ω = Join[ω, {RandomReal[NormalDistribution[μ, σ]]}]
n = 8;
t = t + 1;
o = o + 1;
```

{10.0118, 9.82632, 10.0479, 9.99313, 10.2387, 9.95512, 10.2784, 9.45227}

```
U8 = Utility[ω]
U8next = NextRoundUtility[μpriorlist[[Length[μpriorlist]]], t, n]
```

9.25343

9.08124

```
Ulist = Join[Ulist, {{t, U8}}];
Unextlist = Join[Unextlist, {{t, U8next}}];
```

```
ListPlot[{Ulist, Unextlist}, PlotRange → {{0, 8.5}, {0, 12}}, AxesLabel → {"Time", "Utility"},
 AxesStyle → Thick, LabelStyle → {Bold, Medium}, Filling → {1 → {{2}, {Red, Blue}}}]
```

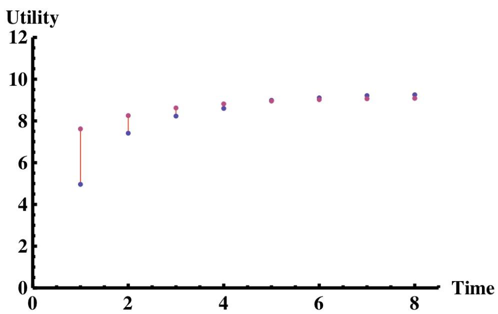


```
ω = Join[ω, {RandomReal[NormalDistribution[μ, σ]]}]
n = 9;
t = t + 1;
o = o + 1;
```

{10.0118, 9.82632, 10.0479, 9.99313, 10.2387, 9.95512, 10.2784, 9.45227, 9.96782}

```
U9 = Utility[ω]
U9next = NextRoundUtility[μpriorlist[[Length[μpriorlist]]], t, n]
```

9.27288

9.0868

```
Ulist = Join[Ulist, {{t, U9}}];
Unextlist = Join[Unextlist, {{t, U9next}}];
```

```
ListPlot[{Ulist, Unextlist}, PlotRange → {{0, 9.5}, {0, 12}}, AxesLabel → {"Time", "Utility"},
 AxesStyle → Thick, LabelStyle → {Bold, Medium}, Filling → {1 → {{2}, {Red, Blue}}}]
```

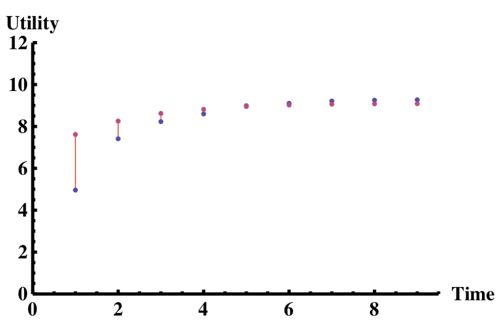

```
ω = Join[ω, {RandomReal[NormalDistribution[μ, σ]]}]
n = 10;
t = t + 1;
o = o + 1;
```

{10.0118, 9.82632, 10.0479, 9.99313, 10.2387, 9.95512, 10.2784, 9.45227, 9.96782, 10.3142}

```
U10 = Utility[ω]
U10next = NextRoundUtility[μpriorlist[[Length[μpriorlist]]], t, n]
```

9.31417

9.08225

```
Ulist = Join[Ulist, {{t, U10}}];
Unextlist = Join[Unextlist, {{t, U10next}}];
```

```
l10 = Line[{{0, U5 - .1}, {10, U5 - .1}}];
Show[ListPlot[{Ulist, Unextlist}, PlotRange → {{0, 10.5}, {0, 12}},
  AxesLabel → {"Time", "Utility"}, AxesStyle → Thick, LabelStyle → {Bold, Medium},
  Filling → {1 → {{2}, {Red, Blue}}}], Graphics[{Dashed, l10}]]
```