Learning to be Risk Averse?

Robert Marks
Economics/AGSM
ASB
UNSW

26th Australasian Economic Theory Workshop
Bond University
7 February 2008

bobm@agsm.edu.au

Learning to be Risk Averse?

Roadmap.

Abstract:

The purpose of this research is to search for the best (highest performing) risk profile of agents who successively choose among risky prospects. An agent's risk profile is his attitude to perceived risk, which can vary from risk preferring to risk neutral (an expected-value decision maker) to risk averse.

Abstract:

The purpose of this research is to search for the best (highest performing) risk profile of agents who successively choose among risky prospects. An agent's risk profile is his attitude to perceived risk, which can vary from risk preferring to risk neutral (an expected-value decision maker) to risk averse.

We use the Genetic Algorithm to search in the complex stochastic space of repeated lotteries.

Abstract:

The purpose of this research is to search for the best (highest performing) risk profile of agents who successively choose among risky prospects. An agent's risk profile is his attitude to perceived risk, which can vary from risk preferring to risk neutral (an expected-value decision maker) to risk averse.

We use the Genetic Algorithm to search in the complex stochastic space of repeated lotteries. We find that agents with a CARA utility function learn to possess risk-neutral risk profiles.

Abstract:

The purpose of this research is to search for the best (highest performing) risk profile of agents who successively choose among risky prospects. An agent's risk profile is his attitude to perceived risk, which can vary from risk preferring to risk neutral (an expected-value decision maker) to risk averse.

We use the Genetic Algorithm to search in the complex stochastic space of repeated lotteries. We find that agents with a CARA utility function learn to possess risk-neutral risk profiles. Since CARA utility functions are wealth-independent, this is not surprising.

Abstract:

The purpose of this research is to search for the best (highest performing) risk profile of agents who successively choose among risky prospects. An agent's risk profile is his attitude to perceived risk, which can vary from risk preferring to risk neutral (an expected-value decision maker) to risk averse.

We use the Genetic Algorithm to search in the complex stochastic space of repeated lotteries. We find that agents with a CARA utility function learn to possess risk-neutral risk profiles. Since CARA utility functions are wealth-independent, this is not surprising. When agents have wealth-dependent, CRRA utility functions, however, they also learn to possess risk profiles that are about risk neutral (from slightly risk-averse to even slightly risk-preferring), which is surprising.

1. Introduction

Informally, it is widely held that in an uncertain world, with the possibility of the discontinuity of bankruptcy, the most prudent risk profile is risk aversion. Indeed, "Risk aversion is one of the most basic assumptions underlying economic behavior" (Szpiro 1997), perhaps because "a dollar that helps us avoid poverty is more valuable than a dollar that helps us become very rich" (Rabin 2000). But is risk aversion the best risk profile? Even with bankruptcy as a possibility?

1. Introduction

Informally, it is widely held that in an uncertain world, with the possibility of the discontinuity of bankruptcy, the most prudent risk profile is risk aversion. Indeed, “Risk aversion is one of the most basic assumptions underlying economic behavior” (Szpiro 1997), perhaps because “a dollar that helps us avoid poverty is more valuable than a dollar that helps us become very rich” (Rabin 2000). But is risk aversion the best risk profile? Even with bankruptcy as a possibility?

To answer this question, we use two kinds of utility function (the wealth-independent exponential utility function, or Constant Absolute Risk Aversion CARA, and the Constant Relative Risk Aversion CRRA function, which is sensitive to the agent’s level of wealth) and run computer experiments in which each agent chooses among three lotteries, and is then awarded with the outcome of the chosen lottery.

Repetition of this choice by many agents allows us use a technique from machine learning — the Genetic Algorithm (Holland 1992) — to search for the best risk profile, where "best" means the highest average payoff when chosing among lotteries.

Repetition of this choice by many agents allows us use a technique from machine learning — the Genetic Algorithm (Holland 1992) — to search for the best risk profile, where "best" means the highest average payoff when chosing among lotteries.

Modelling the agent's utility directly allows us to avoid the indirect inference of Szpiro (1997), who argues that the evolutionary learning technique of the GA does two things: it allows wealth-maximizing agents to succeed even is highly stochastic environments, and it allows the emergence of risk aversion. Indeed, Szpiro argues that risk aversion is the best risk profile to adopt in such an environment.

2. Decisions under Uncertainty and Risk Profiles

The von Neumann-Morgenstern formulation of the decision-maker's attitude to risk is based on the observation that individuals are not always expected-value decision makers. That is, there are situations in which people apparently prefer a lower certain outcome to the higher expected (or probility-weighted) outcome of an uncertain prospect (where the possible outcomes and their possibly subjective, or Bayesian, probabilities are known).

2. Decisions under Uncertainty and Risk Profiles

The von Neumann-Morgenstern formulation of the decision-maker's attitude to risk is based on the observation that individuals are not always expected-value decision makers. That is, there are situations in which people apparently prefer a lower certain outcome to the higher expected (or probility-weighted) outcome of an uncertain prospect (where the possible outcomes and their possibly subjective, or Bayesian, probabilities are known).

An example is paying an insurance premium that is greater than the expected loss without insurance. On the other hand, people will sometimes "gamble" by apparently preferring a lower uncertain outcome to a higher sure thing: this is risk-preferring.

We can formalise this by observing that, by definition, the utility of a lottery is its expected utility, or

$$U(\mathrm{L}) = \sum p_i U(x_i), \qquad (1)$$

where each (discrete) outcome x_i occurs with probability p_i, and $U(x_i)$ is the utility of outcome x_i.

We can formalise this by observing that, by definition, the utility of a lottery is its expected utility, or

$$U(\mathrm{L}) = \sum p_i U(x_i), \tag{1}$$

where each (discrete) outcome x_i occurs with probability p_i, and $U(x_i)$ is the utility of outcome x_i. It is useful to define the Certainty Equivalent $\tilde{x}$ (or C.E.), which is a certain outcome which has the identical utility as the lottery:

$$U(\tilde{x}) \equiv U(\mathrm{L}) = \sum p_i U(x_i). \tag{2}$$

We can formalise this by observing that, by definition, the utility of a lottery is its expected utility, or

$$U(\mathrm{L}) = \sum p_i U(x_i), \tag{1}$$

where each (discrete) outcome x_i occurs with probability p_i, and $U(x_i)$ is the utility of outcome x_i. It is useful to define the Certainty Equivalent $\tilde{x}$ (or C.E.), which is a certain outcome which has the identical utility as the lottery:

$$U(\tilde{x}) \equiv U(\mathrm{L}) = \sum p_i U(x_i). \tag{2}$$

We can use the C.E. to describe the decision-maker's risk profile (Howard 1968).

We can formalise this by observing that, by definition, the utility of a lottery is its expected utility, or

$$U(\mathrm{L}) = \sum p_i U(x_i), \tag{1}$$

where each (discrete) outcome x_i occurs with probability p_i, and $U(x_i)$ is the utility of outcome x_i. It is useful to define the Certainty Equivalent $\tilde{x}$ (or C.E.), which is a certain outcome which has the identical utility as the lottery:

$$U(\tilde{x}) \equiv U(\mathrm{L}) = \sum p_i U(x_i). \tag{2}$$

We can use the C.E. to describe the decision-maker's risk profile (Howard 1968). Define the Expected Value $\bar{x}$ of the Lottery as:

$$\bar{x} = \sum p_i x_i. \tag{3}$$

We can formalise this by observing that, by definition, the utility of a lottery is its expected utility, or

$$U(\mathrm{L}) = \sum p_i U(x_i), \tag{1}$$

where each (discrete) outcome x_i occurs with probability p_i, and $U(x_i)$ is the utility of outcome x_i. It is useful to define the Certainty Equivalent $\tilde{x}$ (or C.E.), which is a certain outcome which has the identical utility as the lottery:

$$U(\tilde{x}) \equiv U(\mathrm{L}) = \sum p_i U(x_i). \tag{2}$$

We can use the C.E. to describe the decision-maker's risk profile (Howard 1968). Define the Expected Value $\bar{x}$ of the Lottery as:

$$\bar{x} = \sum p_i x_i. \tag{3}$$

When $\tilde{x} = \bar{x}$ then the decision-maker's utility function exhibits risk neutrality,

We can formalise this by observing that, by definition, the utility of a lottery is its expected utility, or

$$U(\mathrm{L}) = \sum p_i U(x_i), \tag{1}$$

where each (discrete) outcome x_i occurs with probability p_i, and $U(x_i)$ is the utility of outcome x_i. It is useful to define the Certainty Equivalent $\tilde{x}$ (or C.E.), which is a certain outcome which has the identical utility as the lottery:

$$U(\tilde{x}) \equiv U(\mathrm{L}) = \sum p_i U(x_i). \tag{2}$$

We can use the C.E. to describe the decision-maker's risk profile (Howard 1968). Define the Expected Value $\bar{x}$ of the Lottery as:

$$\bar{x} = \sum p_i x_i. \tag{3}$$

When $\tilde{x} = \bar{x}$ then the decision-maker's utility function exhibits risk neutrality, when $\tilde{x} < \bar{x}$ then risk aversion,

We can formalise this by observing that, by definition, the utility of a lottery is its expected utility, or

$$U(\mathrm{L}) = \sum p_i U(x_i), \tag{1}$$

where each (discrete) outcome x_i occurs with probability p_i, and $U(x_i)$ is the utility of outcome x_i. It is useful to define the Certainty Equivalent $\tilde{x}$ (or C.E.), which is a certain outcome which has the identical utility as the lottery:

$$U(\tilde{x}) \equiv U(\mathrm{L}) = \sum p_i U(x_i). \tag{2}$$

We can use the C.E. to describe the decision-maker's risk profile (Howard 1968). Define the Expected Value $\bar{x}$ of the Lottery as:

$$\bar{x} = \sum p_i x_i. \tag{3}$$

When $\tilde{x} = \bar{x}$ then the decision-maker's utility function exhibits risk neutrality, when $\tilde{x} < \bar{x}$ then risk aversion, and when $\tilde{x} > \bar{x}$ then risk preference.

Approximating the Certainty Equivalent

Expand utility $U(.)$ about the expected value $\bar{x}$.

$$U(x_0) \approx U(\bar{x}) + (x_0 - \bar{x})U'(\bar{x}) + \tfrac{1}{2}(x_o - \bar{x})^2 U''(\bar{x})$$

Approximating the Certainty Equivalent

Expand utility $U(.)$ about the expected value $\bar{x}$.

$$U(x_0) \approx U(\bar{x}) + (x_0 - \bar{x})U'(\bar{x}) + \tfrac{1}{2}(x_o - \bar{x})^2 U''(\bar{x})$$

The C. E. $\tilde{x}$ of a continuous lottery is obtained by integration over the probability density function (p.d.f.) $f_x(.)$:

$$\begin{aligned} U(\tilde{x}) &= \int dx_0 U(x_0) f_x(x_0) \\ &\approx U(\bar{x}) + 0 + \tfrac{1}{2}\sigma^2 U''(\bar{x}) \end{aligned} \qquad (4)$$

Approximating the Certainty Equivalent

Expand utility $U(.)$ about the expected value $\bar{x}$.

$$U(x_0) \approx U(\bar{x}) + (x_0 - \bar{x})U'(\bar{x}) + \tfrac{1}{2}(x_o - \bar{x})^2 U''(\bar{x})$$

The C. E. $\tilde{x}$ of a continuous lottery is obtained by integration over the probability density function (p.d.f.) $f_x(.)$:

$$\begin{aligned} U(\tilde{x}) &= \int dx_0 U(x_0) f_x(x_0) \\ &\approx U(\bar{x}) + 0 + \tfrac{1}{2}\sigma^2 U''(\bar{x}) \end{aligned} \qquad (4)$$

But, by expansion,

$$U(\tilde{x}) \approx U(\bar{x}) + (\tilde{x} - \bar{x})U'(\bar{x}). \qquad (5)$$

Approximating the Certainty Equivalent

Expand utility $U(.)$ about the expected value $\bar{x}$.

$$U(x_0) \approx U(\bar{x}) + (x_0 - \bar{x})U'(\bar{x}) + \tfrac{1}{2}(x_o - \bar{x})^2 U''(\bar{x})$$

The C. E. $\tilde{x}$ of a continuous lottery is obtained by integration over the probability density function (p.d.f.) $f_x(.)$:

$$\begin{aligned} U(\tilde{x}) &= \int dx_0 U(x_0) f_x(x_0) \\ &\approx U(\bar{x}) + 0 + \tfrac{1}{2}\sigma^2 U''(\bar{x}) \end{aligned} \qquad (4)$$

But, by expansion,

$$U(\tilde{x}) \approx U(\bar{x}) + (\tilde{x} - \bar{x})U'(\bar{x}). \qquad (5)$$

Therefore, from equations (4) and (5),

$$\begin{aligned} \tilde{x} - \bar{x} &\approx \tfrac{1}{2}\sigma^2 \frac{U''(\bar{x})}{U'(\bar{x})} \\ \therefore\ \tilde{x} &\approx \bar{x} + \tfrac{1}{2}\sigma^2 \frac{U''(\bar{x})}{U'(\bar{x})} \end{aligned} \qquad (6)$$

Risk aversion

Risk aversion

Risk aversion is not indicated by the slope of the utility curve: it's the *curvature*: if the utility curve is locally —

1. **linear (say, at a point of inflection), then the decision maker is *locally risk neutral.***

2.

Risk aversion

Risk aversion is not indicated by the slope of the utility curve: it's the *curvature*: if the utility curve is locally —

1. **linear (say, at a point of inflection), then the decision maker is *locally risk neutral.***
2. **concave (its slope is decreasing — Diminishing Marginal Utility), then the decision maker is *locally risk averse;***
3.

Risk aversion

Risk aversion is not indicated by the slope of the utility curve: it's the *curvature*: if the utility curve is locally —

1. **linear (say, at a point of inflection), then the decision maker is *locally risk neutral.***
2. **concave (its slope is decreasing — Diminishing Marginal Utility), then the decision maker is *locally risk averse;***
3. **convex (its slope is increasing), then the decision maker is *locally risk preferring.***

Risk-averse, risk-neutral, and risk-preferring exponential utility functions.

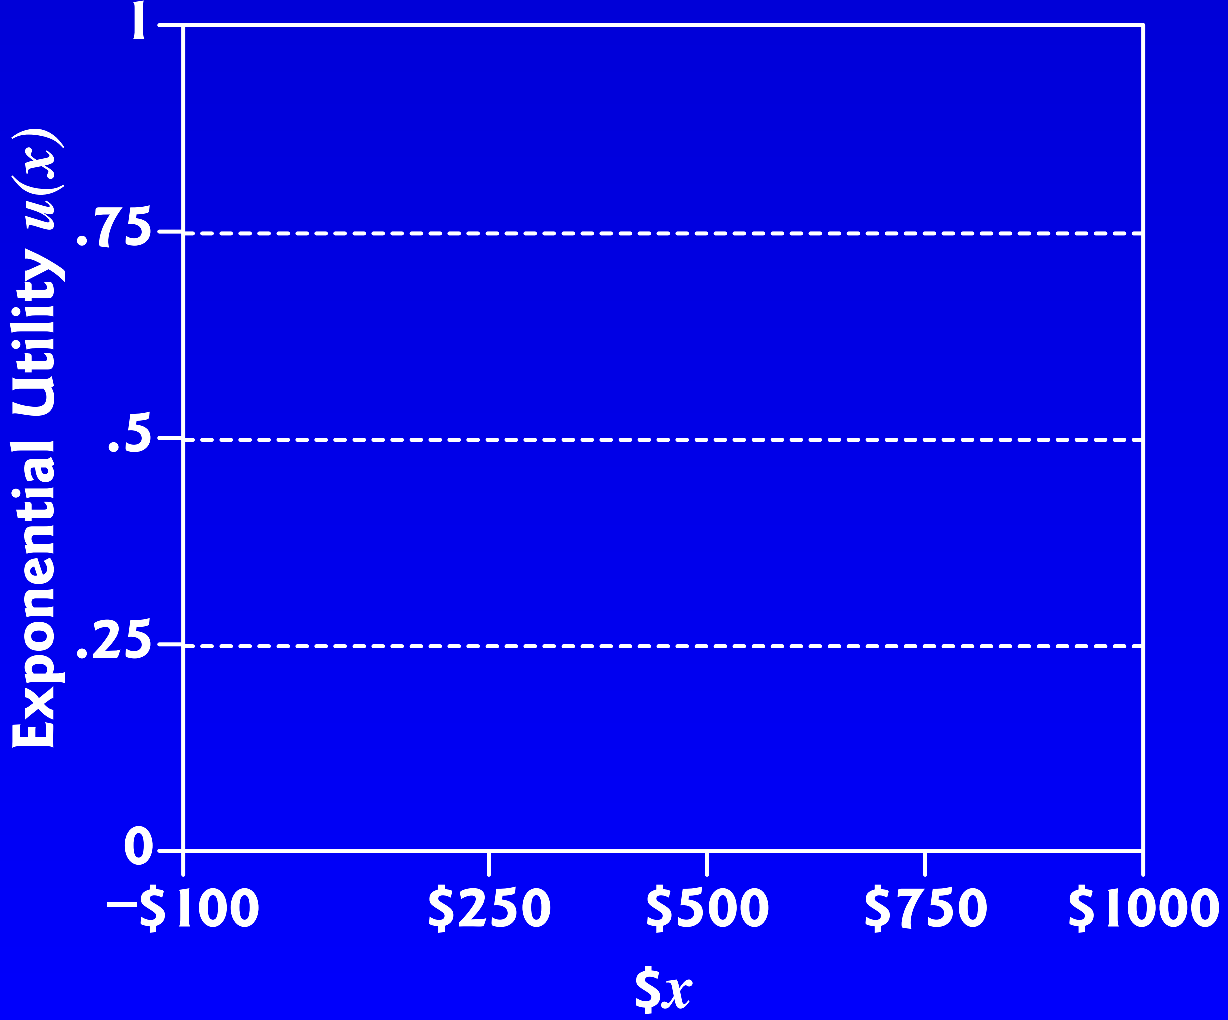

Risk-averse, risk-neutral, and risk-preferring exponential utility functions.

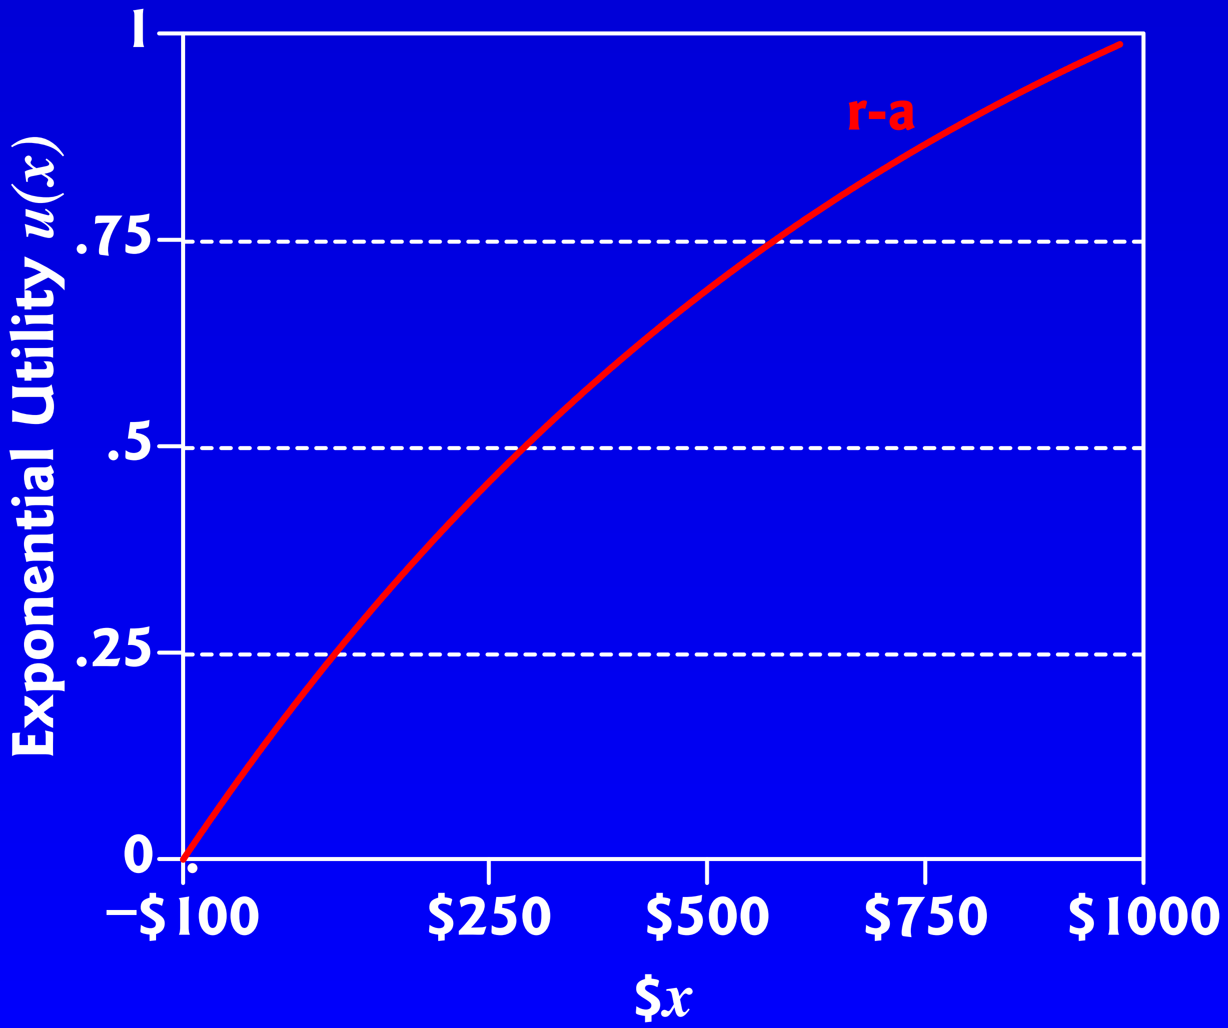

Risk-averse, risk-neutral, and risk-preferring exponential utility functions.

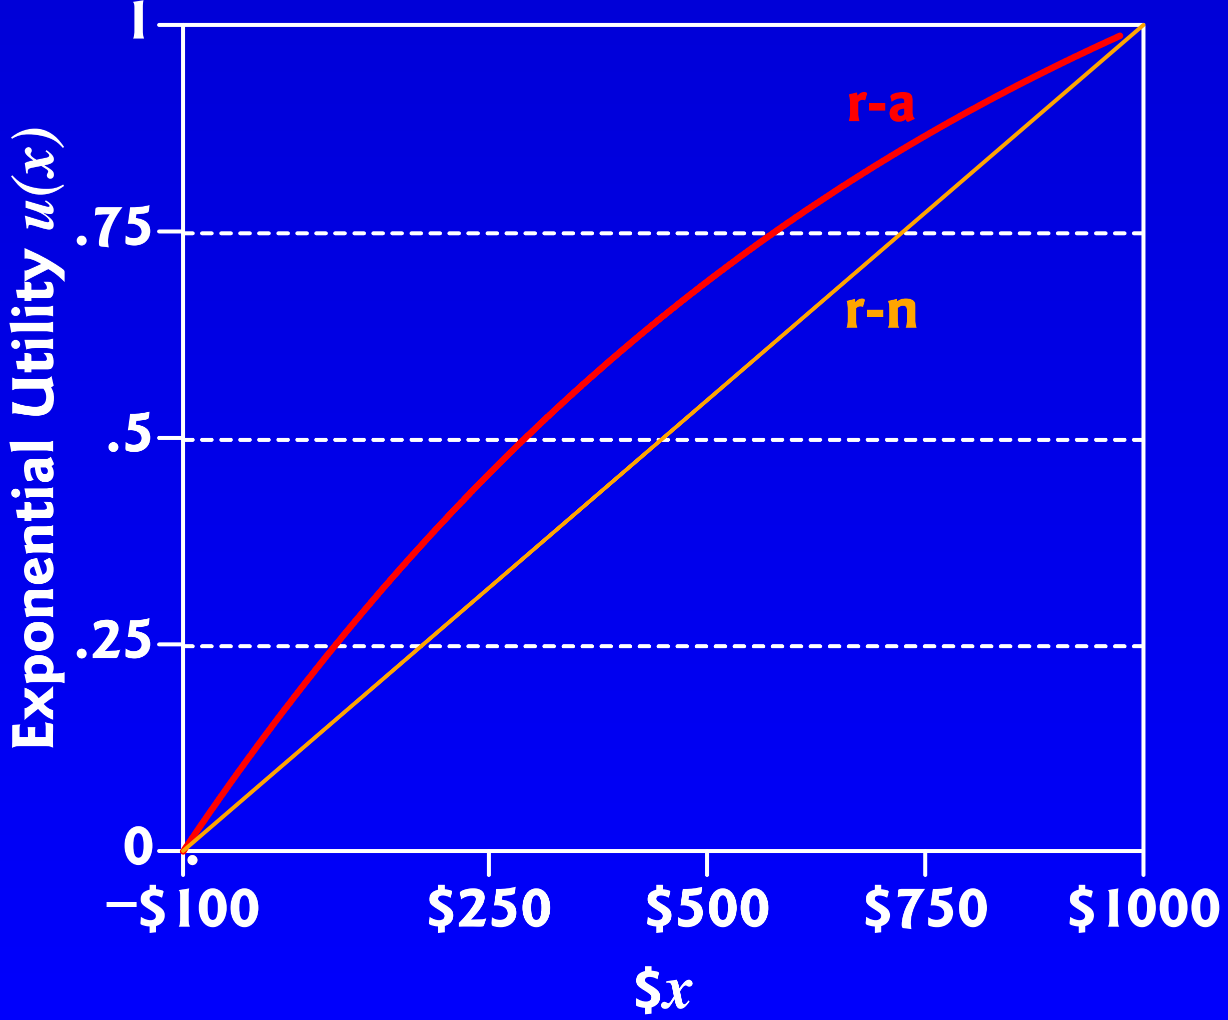

Risk-averse, risk-neutral, and risk-preferring exponential utility functions.

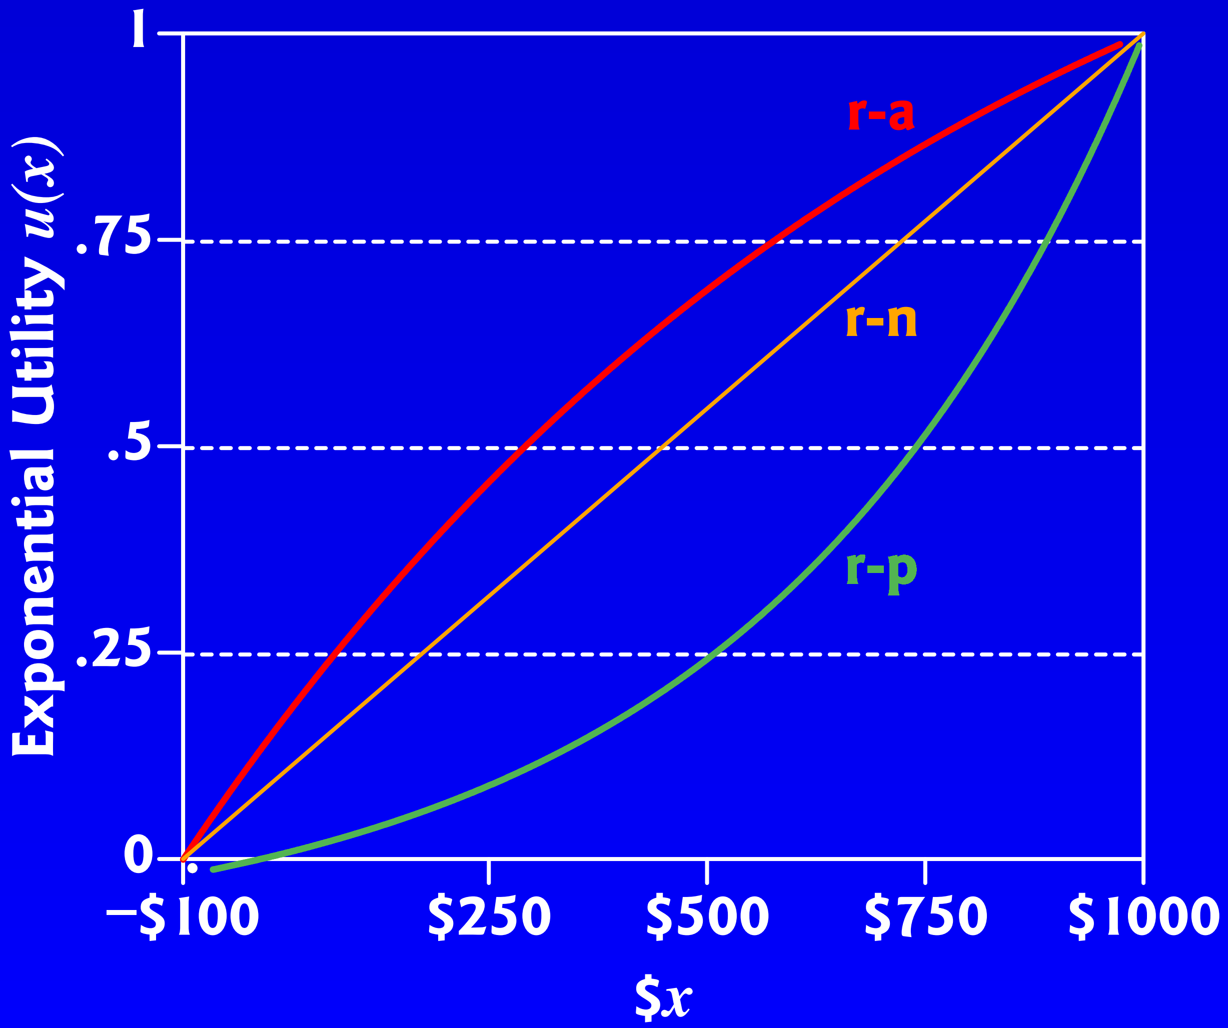

3. Utility Functions

We consider two types of utility function: those which exhibit constant risk preference across all outcomes (so-called wealth-independent utility functions, or Constant Absolute Risk Aversion CARA functions), and those where the risk preference is a function of the wealth of the decision maker (the Constant Relative Risk Aversion CRRA functions).

Wealth Independence

If an increase of all outcomes in a lottery by an amount Δ increases the C.E. by Δ, then the decision maker exhibits wealth independence:

$$U(\tilde{x} + \Delta) = U(\mathrm{L}') = \sum p_i U(x_i + \Delta).$$

Wealth Independence

If an increase of all outcomes in a lottery by an amount Δ increases the C.E. by Δ, then the decision maker exhibits wealth independence:

$$U(\tilde{x} + \Delta) = U(\mathrm{L}') = \sum p_i U(x_i + \Delta).$$

Acceptance of this property restricts possible utility functions to be linear (risk neutral) or exponential, the *wealth-independent,* or constant-absolute-risk-aversion (CARA) functions.

Wealth Independence

If an increase of all outcomes in a lottery by an amount Δ increases the C.E. by Δ, then the decision maker exhibits wealth independence:

$$U(\tilde{x} + \Delta) = U(\mathrm{L}') = \sum p_i U(x_i + \Delta).$$

Acceptance of this property restricts possible utility functions to be linear (risk neutral) or exponential, the *wealth-independent,* or constant-absolute-risk-aversion (CARA) functions.

Acceptance of wealth independence leads to the characterisation of risk preference by a single number, the *risk aversion coefficient,* γ.

Wealth Independence

If an increase of all outcomes in a lottery by an amount Δ increases the C.E. by Δ, then the decision maker exhibits wealth independence:

$$U(\tilde{x} + \Delta) = U(\mathrm{L}') = \sum p_i U(x_i + \Delta).$$

Acceptance of this property restricts possible utility functions to be linear (risk neutral) or exponential, the *wealth-independent,* or constant-absolute-risk-aversion (CARA) functions.

Acceptance of wealth independence leads to the characterisation of risk preference by a single number, the *risk aversion coefficient,* γ.

Since CARA utility functions are wealth-independent, any aversion to bankruptcy is thus precluded, by definition.

3.1 CARA Utility Functions

When utility is linear in outcomes, the decision maker is risk-neutral, across all outcomes, but such a simple constant-risk-profile utility function is of no further interest.

3.1 CARA Utility Functions

When utility is linear in outcomes, the decision maker is risk-neutral, across all outcomes, but such a simple constant-risk-profile utility function is of no further interest. Instead, we consider the exponential CARA functions, where utility U is given by

$$U(x) = 1 - e^{-\gamma x}, \qquad (7)$$

where $U(0) = 0$ and $U(\infty) = 1$, and

3.1 CARA Utility Functions

When utility is linear in outcomes, the decision maker is risk-neutral, across all outcomes, but such a simple constant-risk-profile utility function is of no further interest. Instead, we consider the exponential CARA functions, where utility *U* is given by

$$U(x) = 1 - e^{-\gamma x}, \tag{7}$$

**where $U(0) = 0$ and $U(\infty) = 1$, and
where γ is the *risk aversion coefficient*:**

$$\gamma \equiv -\frac{U''(x)}{U'(x)}. \tag{8}$$

Risk Aversion with Exponential Utility

Risk Aversion with Exponential Utility

From equations (6) and (8), for exponential utility,

$$\tilde{x} \approx \bar{x} - \tfrac{1}{2}\sigma^2\gamma$$

which indicates that when $\gamma = 0$, then $\tilde{x} \approx \bar{x}$ (risk neutrality), when $\gamma > 0$, then $\tilde{x} < \bar{x}$ (risk averse), and when $\gamma < 0$, then $\tilde{x} > \bar{x}$ (risk preferring), with positive variance.

Risk Aversion with Exponential Utility

From equations (6) and (8), for exponential utility,

$$\tilde{x} \approx \bar{x} - \tfrac{1}{2}\sigma^2\gamma$$

which indicates that when $\gamma = 0$, then $\tilde{x} \approx \bar{x}$ (risk neutrality), when $\gamma > 0$, then $\tilde{x} < \bar{x}$ (risk averse), and when $\gamma < 0$, then $\tilde{x} > \bar{x}$ (risk preferring), with positive variance.

Summarizing this:

Sign of γ	Risk profile	Curvature
$\gamma = 0$	risk neutral	$U''(x) = 0$
$\gamma > 0$	risk averse	$U''(x) < 0$
$\gamma < 0$	risk preferring	$U''(x) > 0$

3.2 CRRA Utility Functions

We want a utility function which is *not* wealth-independent, to see whether that will result in risk-averse agents doing best.

3.2 CRRA Utility Functions

We want a utility function which is *not* wealth-independent, to see whether that will result in risk-averse agents doing best.

The Arrow-Pratt measure of relative risk aversion (RRA) is defined as

$$\rho(w) = -w \frac{U''(w)}{U'(w)} = w\gamma \tag{9}$$

This introduces wealth w into the agent's risk preferences, so that lower wealth can be associated with higher risk aversion. Risk aversion coefficient γ is as in equation (8).

3.2 CRRA Utility Functions

We want a utility function which is *not* wealth-independent, to see whether that will result in risk-averse agents doing best.

The Arrow-Pratt measure of relative risk aversion (RRA) is defined as

$$\rho(w) = -w\,\frac{U''(w)}{U'(w)} = w\gamma \tag{9}$$

This introduces wealth w into the agent's risk preferences, so that lower wealth can be associated with higher risk aversion. Risk aversion coefficient γ is as in equation (8).

The Constant Elasticity of Substitution (CES) utility function:

$$U(w) = \frac{w^{1-\rho}}{1-\rho}, \quad w > 0, \tag{10}$$

exhibits CRRA, equation (9).

Risk Aversion with CES Utility

In the CRRA simulations, we use the cumulative sum of the realisations of payoffs won in previous lotteries chosen by the agent plus the possible payoff in this lottery as the wealth w in equation (10).

Risk Aversion with CES Utility

In the CRRA simulations, we use the cumulative sum of the realisations of payoffs won in previous lotteries chosen by the agent plus the possible payoff in this lottery as the wealth w in equation (10).

Each agent codes for ρ.

Risk Aversion with CES Utility

In the CRRA simulations, we use the cumulative sum of the realisations of payoffs won in previous lotteries chosen by the agent plus the possible payoff in this lottery as the wealth w in equation (10).

Each agent codes for ρ.

From equation (6), the C.E. with CES utility is approximated by

$$\tilde{x} \approx \bar{x} - \frac{1}{2} \frac{\rho}{w} \sigma^2$$

Risk Aversion with CES Utility

In the CRRA simulations, we use the cumulative sum of the realisations of payoffs won in previous lotteries chosen by the agent plus the possible payoff in this lottery as the wealth w in equation (10).

Each agent codes for ρ.

From equation (6), the C.E. with CES utility is approximated by

$$\tilde{x} \approx \bar{x} - \frac{1}{2}\,\frac{\rho}{w}\,\sigma^2$$

Iff $\frac{1}{2}\,\frac{\rho}{w}\,\sigma^2 > 0$ (or $\rho/w > 0$), then C.E. $\tilde{x}$ < expected mean $\bar{x}$, and the decision maker is risk averse.

Risk Aversion with CES Utility

In the CRRA simulations, we use the cumulative sum of the realisations of payoffs won in previous lotteries chosen by the agent plus the possible payoff in this lottery as the wealth w in equation (10).

Each agent codes for ρ.

From equation (6), the C.E. with CES utility is approximated by

$$\tilde{x} \approx \bar{x} - \frac{1}{2}\frac{\rho}{w}\sigma^2$$

Iff $\frac{1}{2}\frac{\rho}{w}\sigma^2 > 0$ (or $\rho/w > 0$), then C.E. $\tilde{x}$ < expected mean $\bar{x}$, and the decision maker is risk averse.

With $w > 0$, $\rho > 0$ is equivalent to risk aversion.

4. The Simulations

Each lottery is randomly constructed: the two payoffs ("prizes") are randomly chosen in the interval from 1 to 2 $\times$ MAP, usually 100; and the probability is also chosen randomly. (Each lottery has, of course, a single degree of freedom for probability). Each agent calculates the expected utility of each of the three lotteries, using its utility function (a function of its γ or ρ/W), and chooses the lottery with the highest expected utility. To do this, agents know the prizes and probabilities of all three lotteries.

4. The Simulations

Each lottery is randomly constructed: the two payoffs (“prizes”) are randomly chosen in the interval from 1 to 2 × MAP, usually 100; and the probability is also chosen randomly. (Each lottery has, of course, a single degree of freedom for probability). Each agent calculates the expected utility of each of the three lotteries, using its utility function (a function of its γ or ρ/W), and chooses the lottery with the highest expected utility. To do this, agents know the prizes and probabilities of all three lotteries.

Then the actual (simulated) outcome of the chosen lottery is randomly realised, using its probability. The winnings of the Constant Absolute Risk Aversion agent (respectively, the wealth of the Constant Relative Risk Aversion agent) is incremented accordingly. Each agent chooses 1000 lotteries.

4. The Simulations

Each lottery is randomly constructed: the two payoffs ("prizes") are randomly chosen in the interval from 1 to 2 $\times$ MAP, usually 100; and the probability is also chosen randomly. (Each lottery has, of course, a single degree of freedom for probability). Each agent calculates the expected utility of each of the three lotteries, using its utility function (a function of its γ or ρ / W), and chooses the lottery with the highest expected utility. To do this, agents know the prizes and probabilities of all three lotteries.

Then the actual (simulated) outcome of the chosen lottery is randomly realised, using its probability. The winnings of the Constant Absolute Risk Aversion agent (respectively, the wealth of the Constant Relative Risk Aversion agent) is incremented accordingly. Each agent chooses 1000 lotteries. At this stage there is a population of agents, each of which has a average winnings or a cumulative level of wealth, based on its risk profile and the successive outcomes of its choices among the lotteries.

Searching with the Genetic Algorithm

We now use an implementation of the Genetic Algorithm (Gilbert 2004) to search for the best risk profile. That is, we select the best-performing agents to be the “parents” of the next generation of agents, which is generated by “crossover” and “mutation” of the chromosomes of the pairs of parents. Each of the new generation of agents chooses the lottery with highest expected utility a thousand times. Again, the best are selected to be the parents of the next generation.

Searching with the Genetic Algorithm

We now use an implementation of the Genetic Algorithm (Gilbert 2004) to search for the best risk profile. That is, we select the best-performing agents to be the “parents” of the next generation of agents, which is generated by “crossover” and “mutation” of the chromosomes of the pairs of parents. Each of the new generation of agents chooses the lottery with highest expected utility a thousand times. Again, the best are selected to be the parents of the next generation.

We use the GA simulation in this search as an empirical alternative to solving for the best (highest performing) risk profile analytically. Note that Rabin (2000) asserts that “theory actually predicts virtual risk neutrality.” We return to this paper in the Discussion below.

The Simulations with CARA Utility

Using NetLogo (Wilensky 1999), we model each agent as a binary string which codes to its risk-aversion coefficient, γ, in the interval $\pm$ 1.048576.

The Simulations with CARA Utility

Using NetLogo (Wilensky 1999), we model each agent as a binary string which codes to its risk-aversion coefficient, γ, in the interval $\pm$ 1.048576.

Each lottery is a two-prize lottery, where each prize is chosen from a uniform distribution, between 1 and 2 $\times$ MAP (maximum absolute prize), where MAP can be set up to 100 by the simulator, and the single probability is chosen randomly from uniform [0,1].

The Simulations with CARA Utility

Using NetLogo (Wilensky 1999), we model each agent as a binary string which codes to its risk-aversion coefficient, γ, in the interval $\pm$ 1.048576.

Each lottery is a two-prize lottery, where each prize is chosen from a uniform distribution, between 1 and 2 $\times$ MAP (maximum absolute prize), where MAP can be set up to 100 by the simulator, and the single probability is chosen randomly from uniform [0,1].

Each agent chooses the lottery with the highest expected utility from equations (1) and (7), based on its value of γ. Then a realised outcome is calculated for that lottery, based on its probability.

The Simulations with CARA Utility

Using NetLogo (Wilensky 1999), we model each agent as a binary string which codes to its risk-aversion coefficient, γ, in the interval $\pm$ 1.048576.

Each lottery is a two-prize lottery, where each prize is chosen from a uniform distribution, between 1 and 2 $\times$ MAP (maximum absolute prize), where MAP can be set up to 100 by the simulator, and the single probability is chosen randomly from uniform [0,1].

Each agent chooses the lottery with the highest expected utility from equations (1) and (7), based on its value of γ. Then a realised outcome is calculated for that lottery, based on its probability.

Each agent faces 1000 lottery choices, and the cumulative winnings that agent's "fitness" for the Genetic Algorithm.

The CARA Results

See `http://www.agsm.edu.au/~bobm/teaching/SimSS/NetLogo4-models/RA-CARA-EU-312p.html` **for a Java aplet and the NetLogo code.**

The windows captured from the NetLogo simulations show three things clearly:

1.

The CARA Results

See `http://www.agsm.edu.au/~bobm/teaching/SimSS/NetLogo4-models/RA-CARA-EU-312p.html` **for a Java aplet and the NetLogo code.**

The windows captured from the NetLogo simulations show three things clearly:

1. **The mean (black) fitness grows quickly to a plateau after 20 generations or so;**

2.

The CARA Results

See `http://www.agsm.edu.au/~bobm/teaching/SimSS/NetLogo4-models/RA-CARA-EU-312p.html` **for a Java aplet and the NetLogo code.**

The windows captured from the NetLogo simulations show three things clearly:

1. **The mean (black) fitness grows quickly to a plateau after 20 generations or so;**
2. **the mean, maximum, and minimum risk-aversion coefficients γ (resp. black, green, red) converge to close to zero (risk neutrality) over the same period, and**
3.

The CARA Results

See `http://www.agsm.edu.au/~bobm/teaching/SimSS/NetLogo4-models/RA-CARA-EU-312p.html` **for a Java aplet and the NetLogo code.**

The windows captured from the NetLogo simulations show three things clearly:

1. **The mean (black) fitness grows quickly to a plateau after 20 generations or so;**
2. **the mean, maximum, and minimum risk-aversion coefficients γ (resp. black, green, red) converge to close to zero (risk neutrality) over the same period, and**
3. **Any γ deviation from zero up (more risk-averse) or down (more risk-preferring) leads to the minimum (red) fitness in that generation collapsing from close to the mean fitness.**

The CARA Results

See `http://www.agsm.edu.au/~bobm/teaching/SimSS/NetLogo4-models/RA-CARA-EU-312p.html` **for a Java aplet and the NetLogo code.**

The windows captured from the NetLogo simulations show three things clearly:

1. **The mean (black) fitness grows quickly to a plateau after 20 generations or so;**
2. **the mean, maximum, and minimum risk-aversion coefficients γ (resp. black, green, red) converge to close to zero (risk neutrality) over the same period, and**
3. **Any γ deviation from zero up (more risk-averse) or down (more risk-preferring) leads to the minimum (red) fitness in that generation collapsing from close to the mean fitness.**

These observations clearly show that CARA agents perform best (in terms of their lottery winnings) who are closest to risk neutral ($\gamma = 0$).

The CARA Results

See http://www.agsm.edu.au/~bobm/teaching/SimSS/NetLogo4-models/RA-CARA-EU-312p.html **for a Java aplet and the NetLogo code.**

The windows captured from the NetLogo simulations show three things clearly:

1. **The mean (black) fitness grows quickly to a plateau after 20 generations or so;**
2. **the mean, maximum, and minimum risk-aversion coefficients γ (resp. black, green, red) converge to close to zero (risk neutrality) over the same period, and**
3. **Any γ deviation from zero up (more risk-averse) or down (more risk-preferring) leads to the minimum (red) fitness in that generation collapsing from close to the mean fitness.**

These observations clearly show that CARA agents perform best (in terms of their lottery winnings) who are closest to risk neutral (γ = 0).

Too risk averse, and they forgo fair lotteries; too risk preferring and they choose too many risky lotteries.

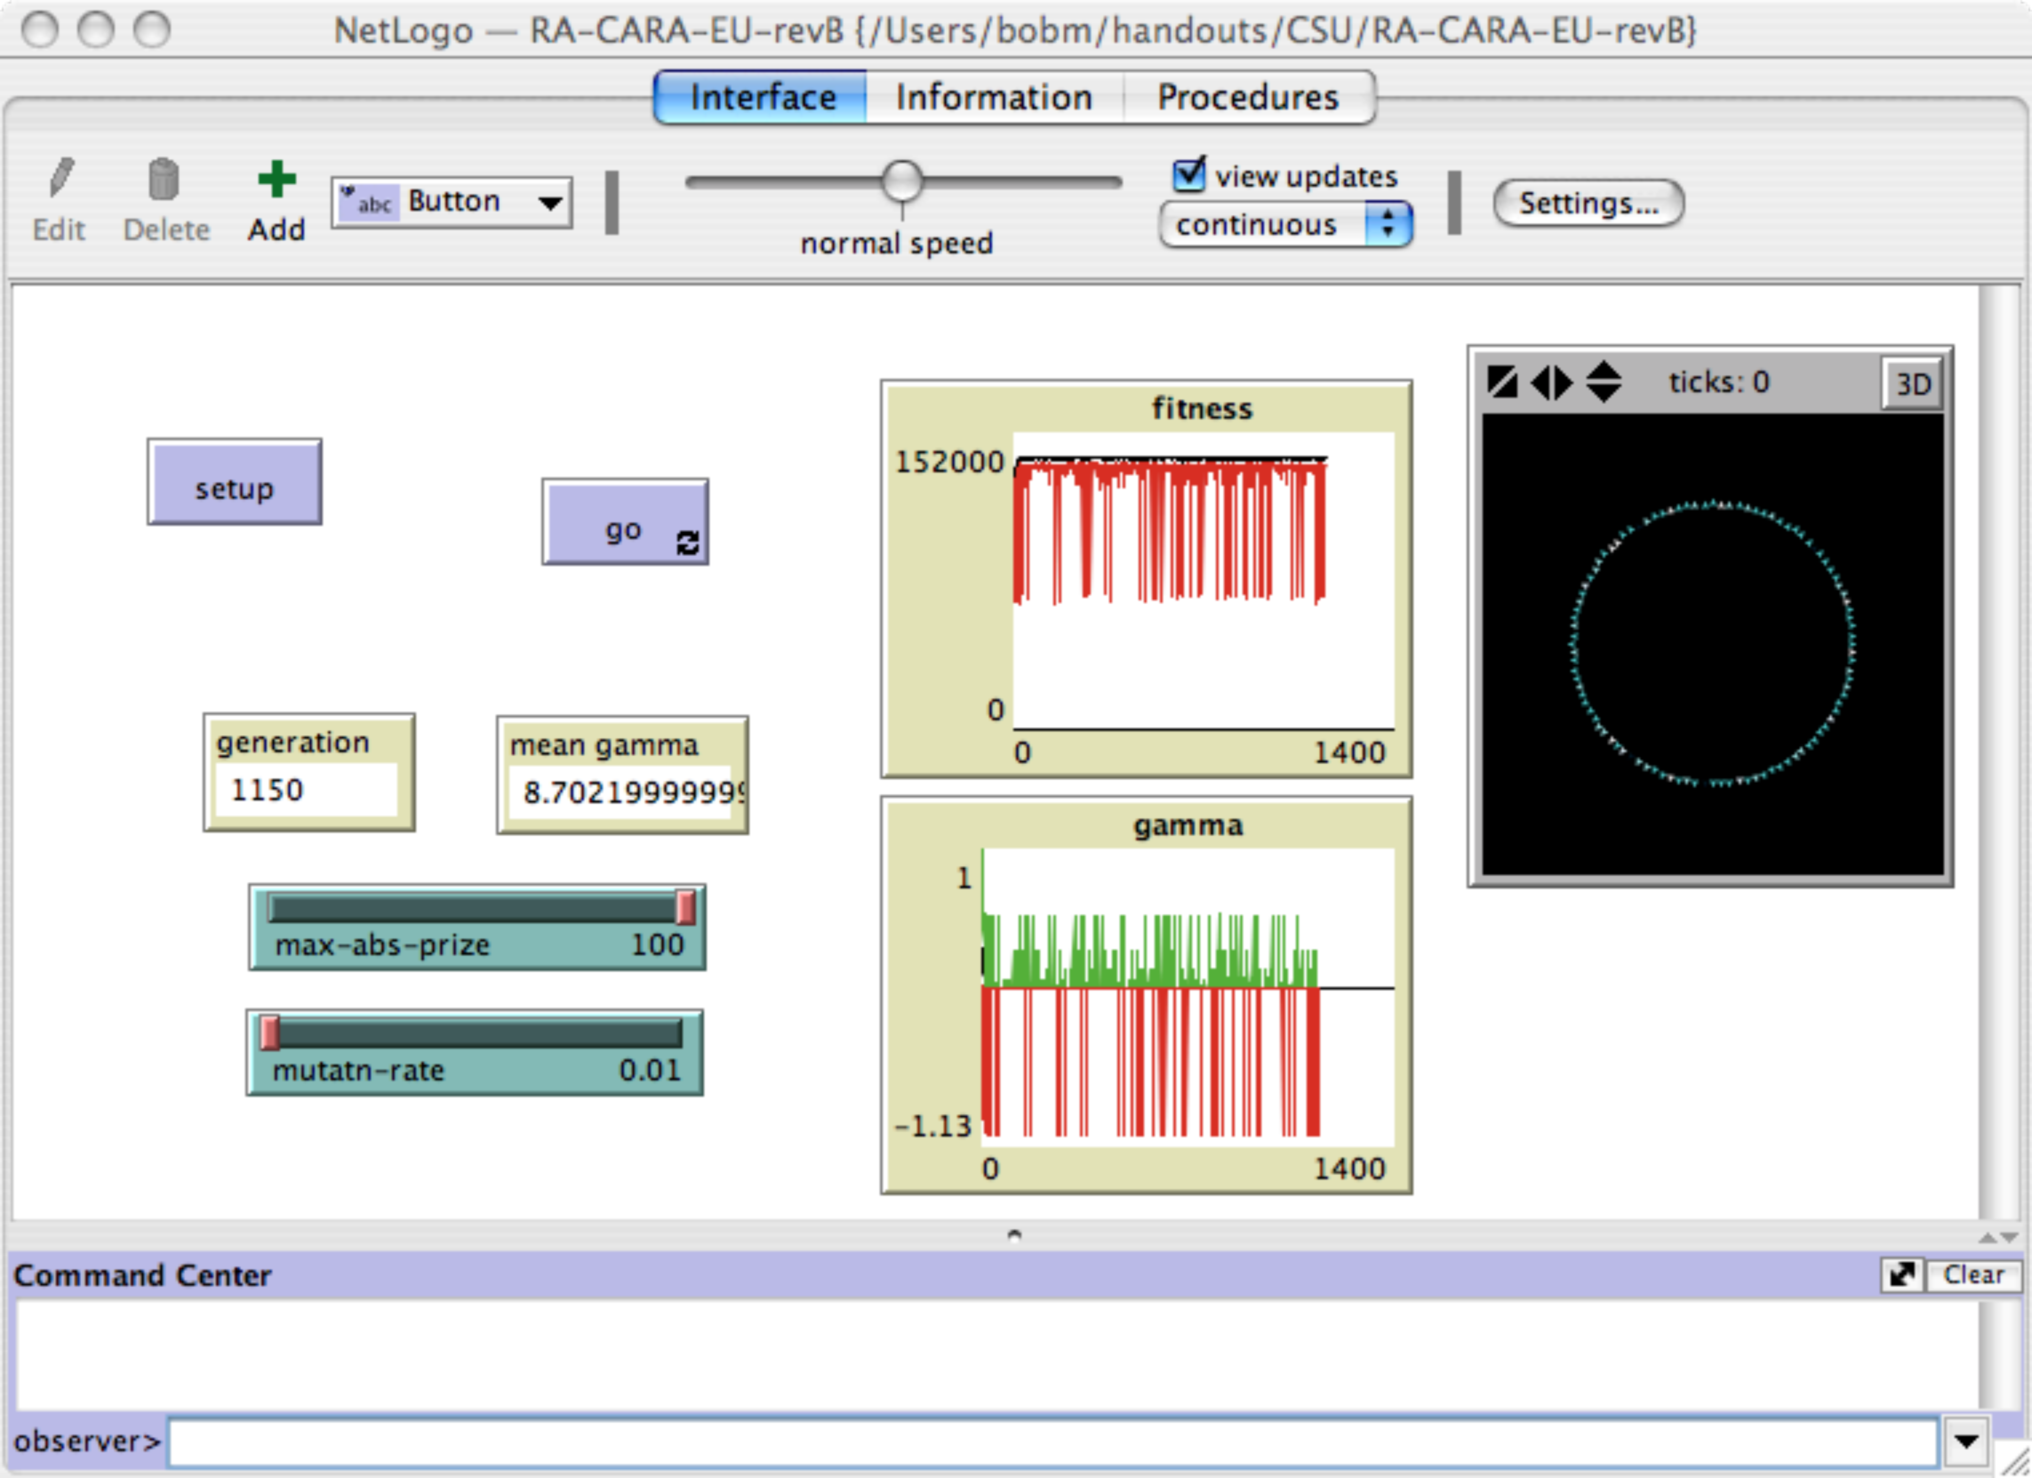
NetLogo — RA-CARA-EU-revB {/Users/bobm/handouts/CSU/RA-CARA-EU-revB}
Interface
Information
Procedures
Edit
Delete
Add
Button
view updates
continuous
normal speed
Settings...
setup
go
generation
1150
mean gamma
8.7021999999
max-abs-prize
100
mutatn-rate
0.01
fitness
152000
0
0
1400
gamma
1
-1.13
0
1400
ticks: 0
3D
Command Center
Clear
observer>

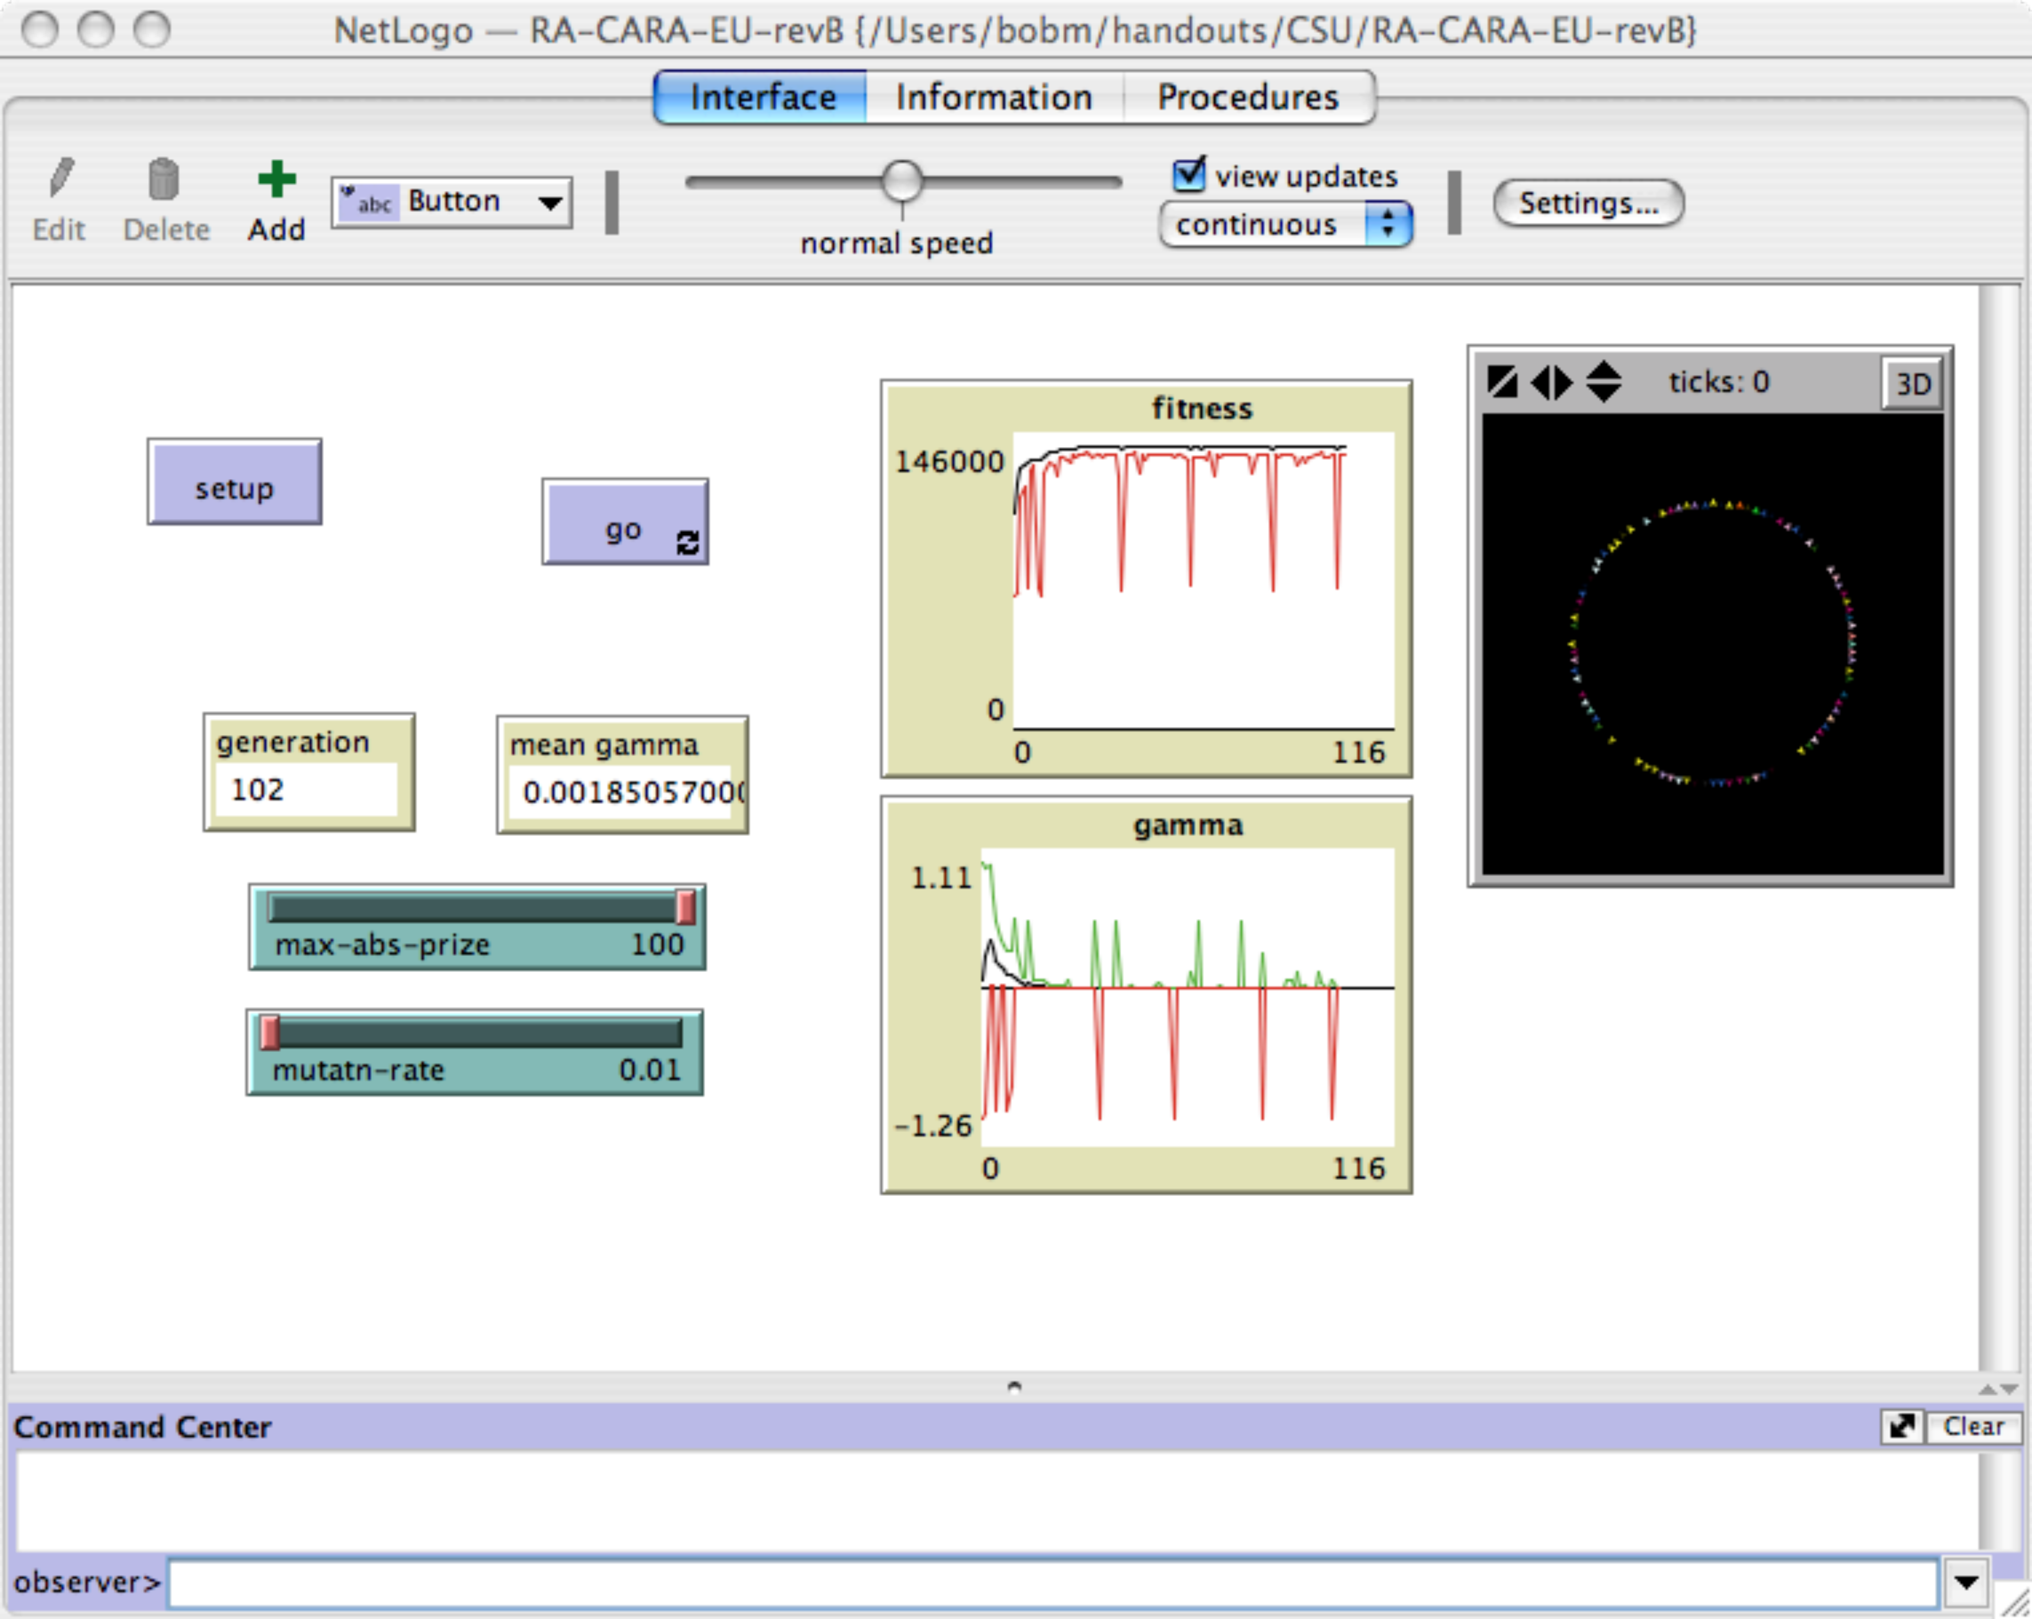

NetLogo — RA-CARA-EU-revB {/Users/bobm/handouts/CSU/RA-CARA-EU-revB}
Interface
Information
Procedures
Edit
Delete
Add
Button
normal speed
view updates
continuous
Settings...
setup
go
generation
102
mean gamma
0.00185057000
max-abs-prize
100
mutatn-rate
0.01
fitness
146000
0
0
116
gamma
1.11
-1.26
0
116
ticks: 0
3D
Command Center
Clear
observer>

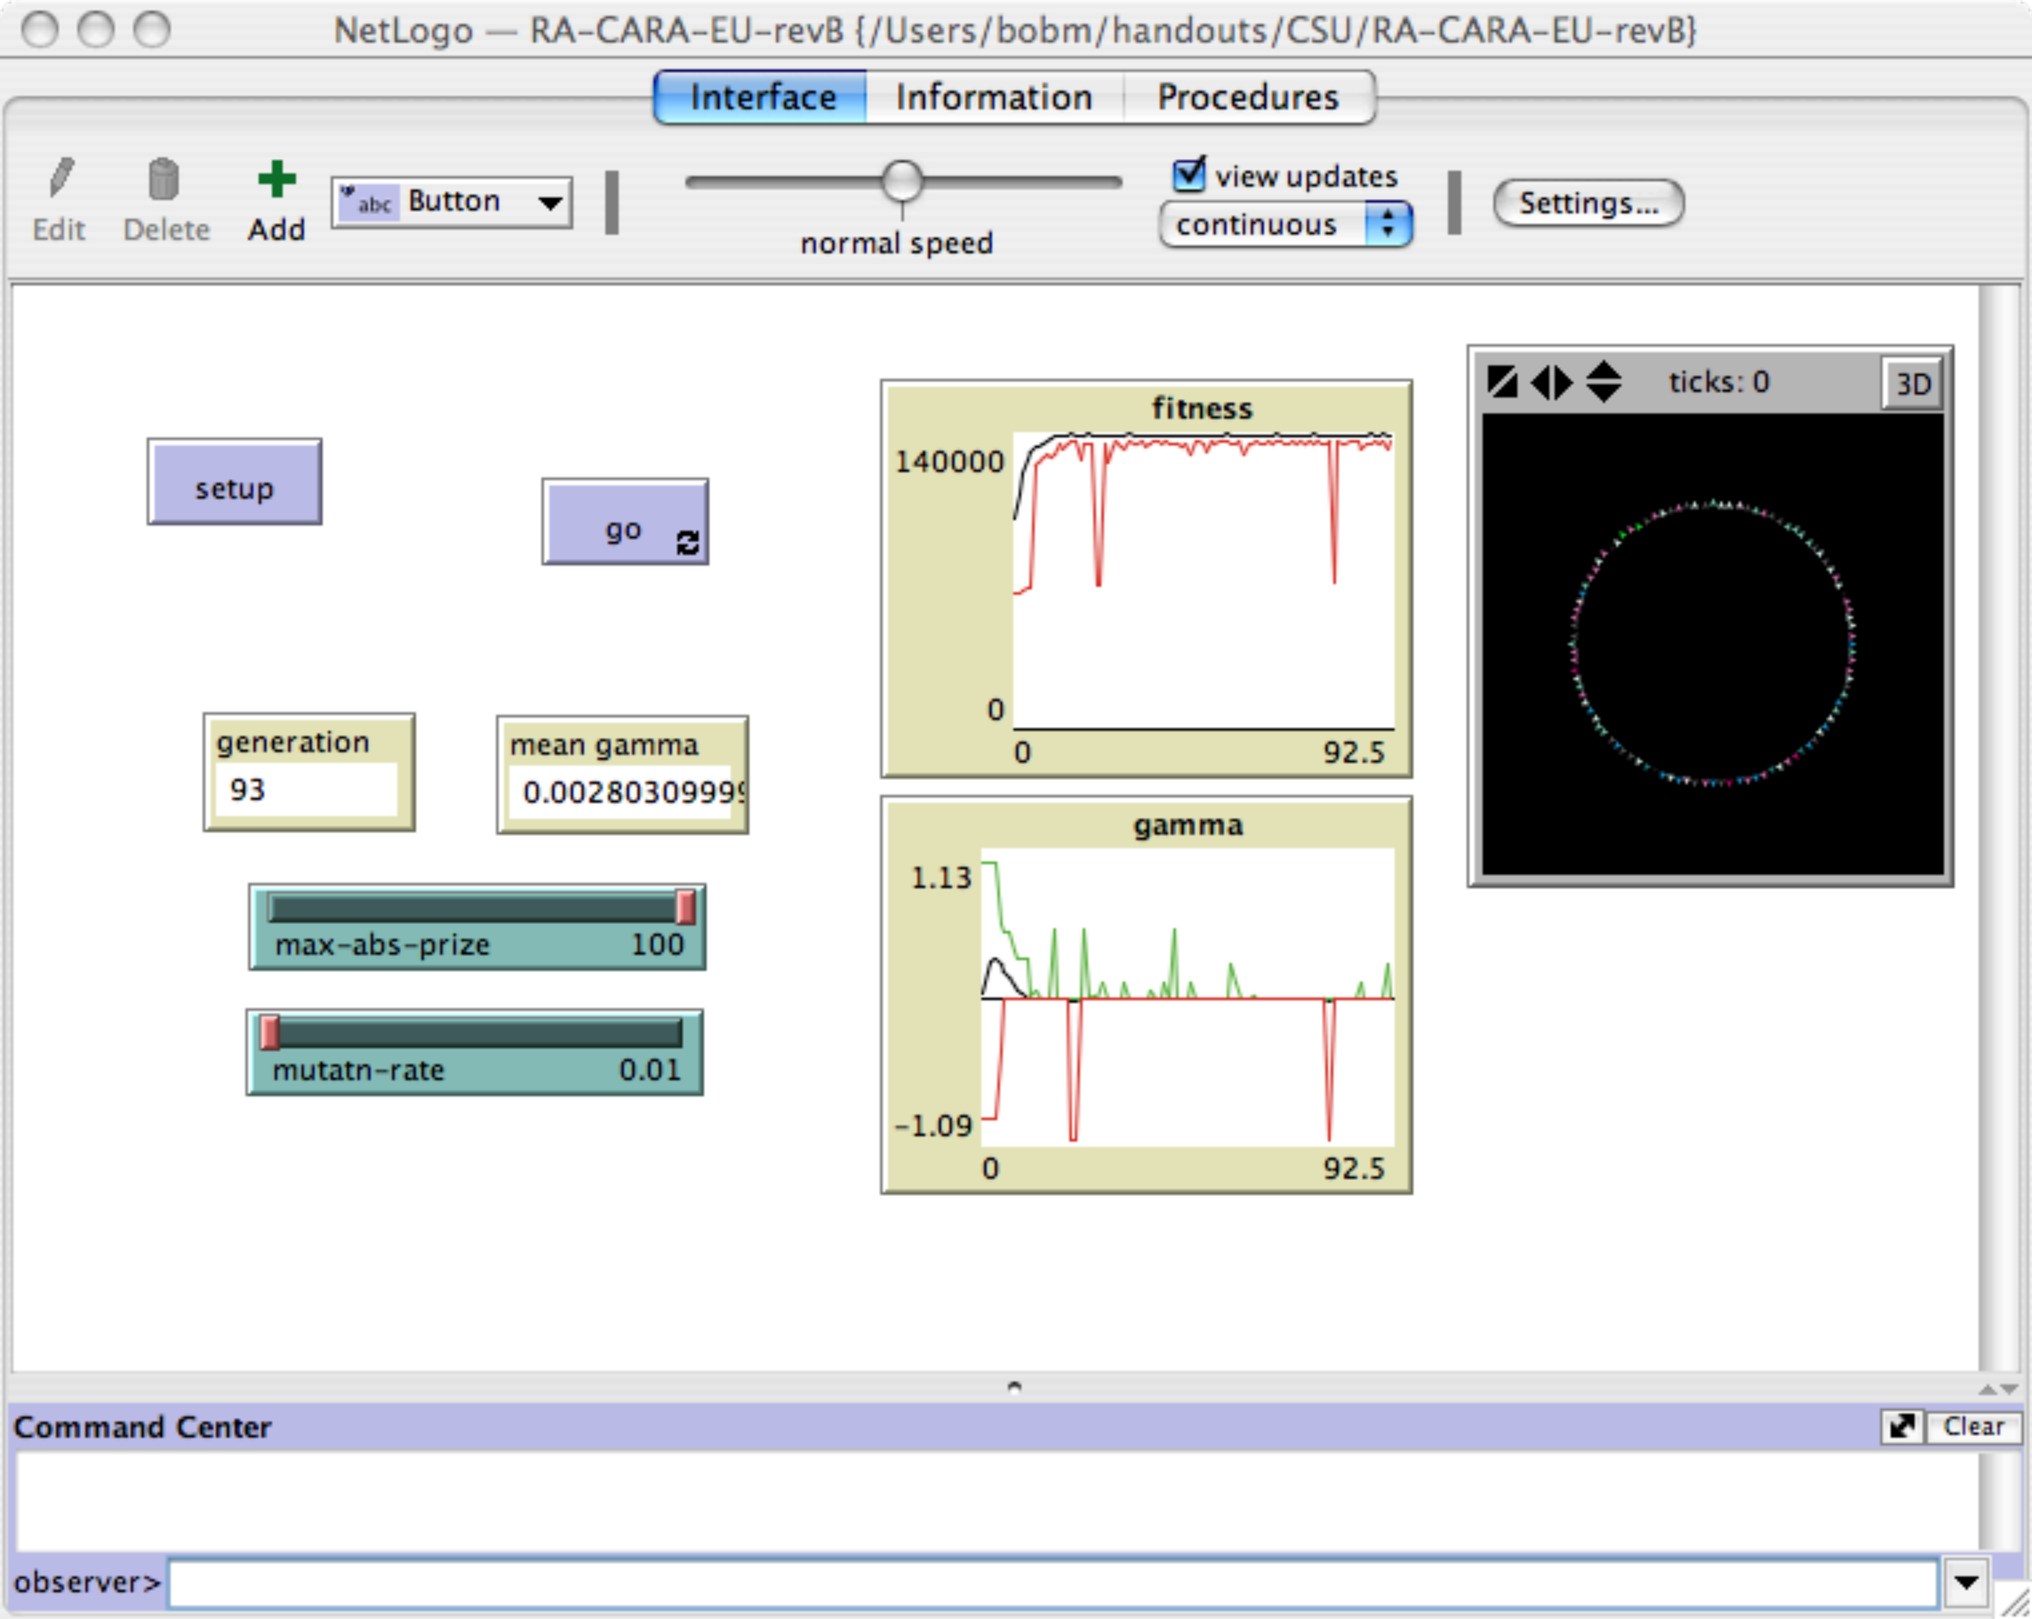
NetLogo — RA-CARA-EU-revB {/Users/bobm/handouts/CSU/RA-CARA-EU-revB}
Interface
Information
Procedures
Edit
Delete
Add
Button
normal speed
view updates
continuous
Settings...
setup
go
generation
93
mean gamma
0.0028030999
max-abs-prize
100
mutatn-rate
0.01
fitness
140000
0
0
92.5
gamma
1.13
-1.09
0
92.5
ticks: 0
3D
Command Center
Clear
observer>

The CRRA Results

See `http://www.agsm.edu.au/~bobm/teaching/SimSS/NetLogo4-models/DRA-CRRA-EU-revCD-312p.html` **for a Java aplet and the NetLogo code.**

The CRRA Results

See `http://www.agsm.edu.au/~bobm/teaching/SimSS/NetLogo4-models/DRA-CRRA-EU-revCD-312p.html` **for a Java aplet and the NetLogo code.**

Despite our prior belief, the CARA agents do not learn to be risk averse, but to be risk neutral. Is this because the wealth-independent CARA utility function precludes bankruptcy?

The CRRA Results

See http://www.agsm.edu.au/~bobm/teaching/SimSS/NetLogo4-models/DRA-CRRA-EU-revCD-312p.html **for a Java aplet and the NetLogo code.**

Despite our prior belief, the CARA agents do not learn to be risk averse, but to be risk neutral. Is this because the wealth-independent CARA utility function precludes bankruptcy?

We could, of course, put a floor on agent wealth, below which is oblivion, but better to use a utility formulation that is not wealth independent and repeat the search.

The CRRA Results

See `http://www.agsm.edu.au/~bobm/teaching/SimSS/NetLogo4-models/DRA-CRRA-EU-revCD-312p.html` **for a Java aplet and the NetLogo code.**

Despite our prior belief, the CARA agents do not learn to be risk averse, but to be risk neutral. Is this because the wealth-independent CARA utility function precludes bankruptcy?

We could, of course, put a floor on agent wealth, below which is oblivion, but better to use a utility formulation that is not wealth independent and repeat the search. We use the CES utility functions (equation (10)) that exhibits CRRA.

The CRRA Results

See `http://www.agsm.edu.au/~bobm/teaching/SimSS/NetLogo4-models/DRA-CRRA-EU-revCD-312p.html` **for a Java aplet and the NetLogo code.**

Despite our prior belief, the CARA agents do not learn to be risk averse, but to be risk neutral. Is this because the wealth-independent CARA utility function precludes bankruptcy?

We could, of course, put a floor on agent wealth, below which is oblivion, but better to use a utility formulation that is not wealth independent and repeat the search. We use the CES utility functions (equation (10)) that exhibits CRRA.

The results are not surprising: the CRRA agents learn to be slightly risk averse (or risk neutral or slightly risk-preferring with $\rho < 0$).

Remember: $\gamma = \dfrac{\rho}{w}$, so dividing the ρ values by the high w values attained implies corresponding minute values of γ here.

“

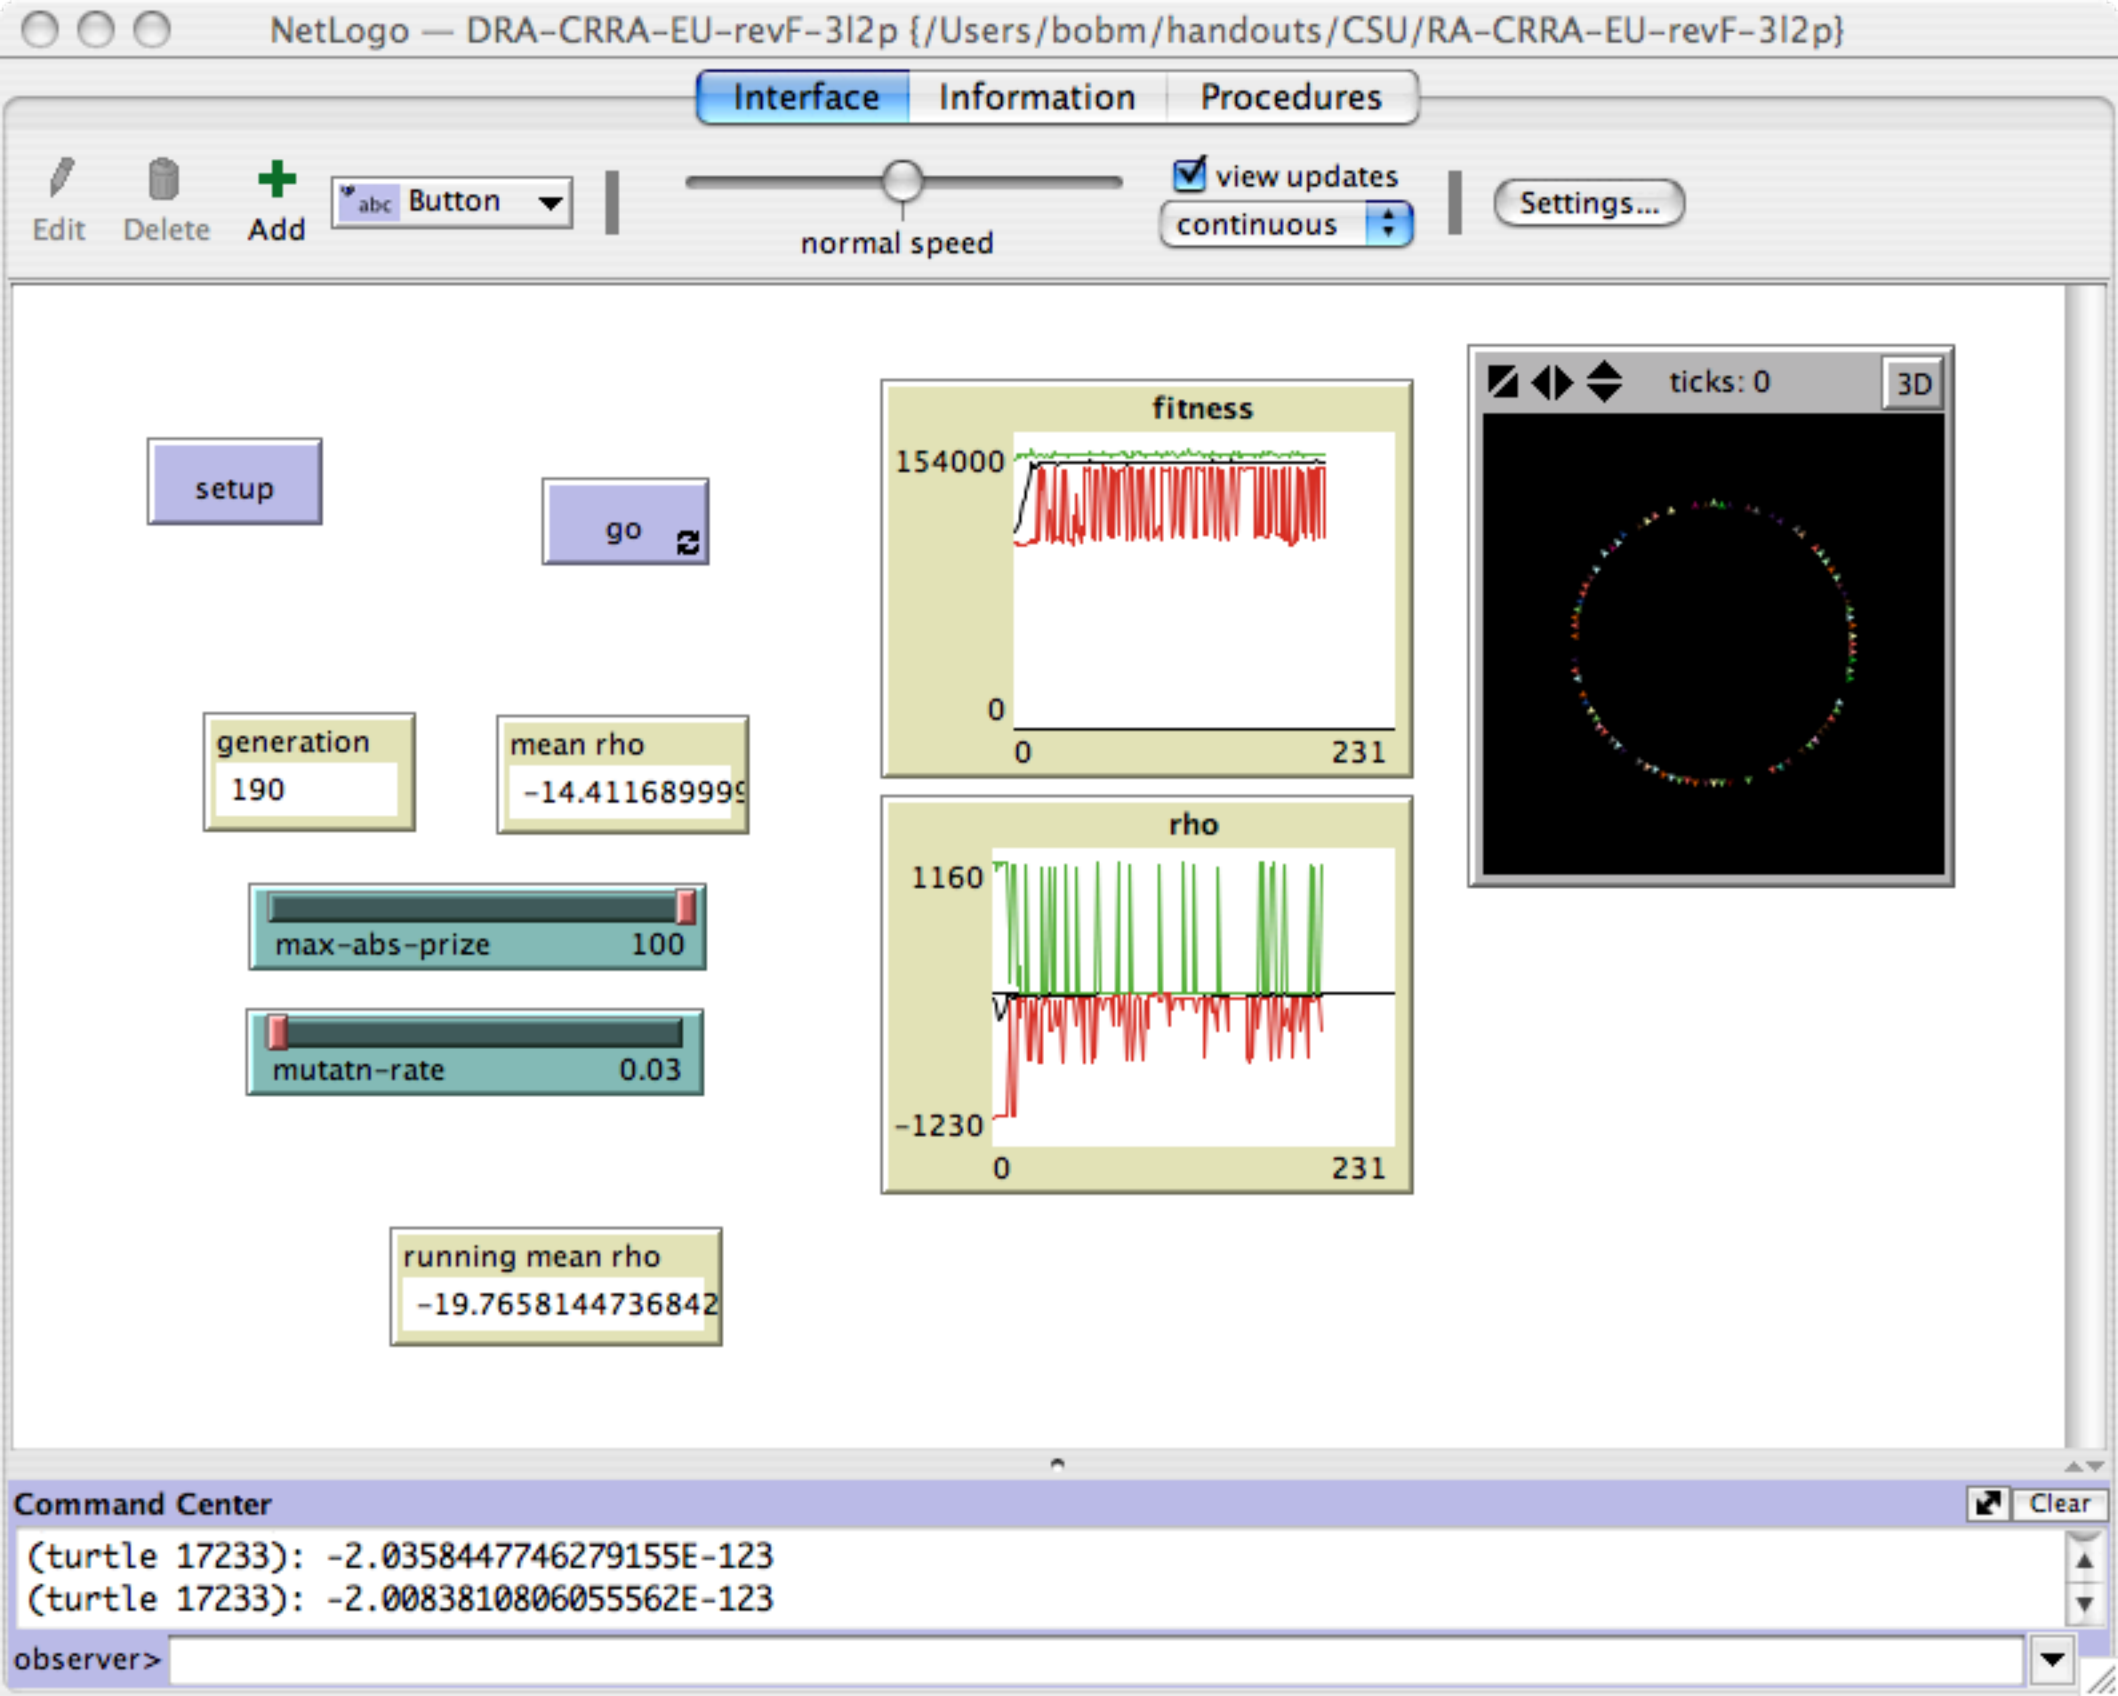
NetLogo — DRA-CRRA-EU-revF-3l2p {/Users/bobm/handouts/CSU/RA-CRRA-EU-revF-3l2p}
Interface
Information
Procedures
Edit
Delete
Add
Button
normal speed
view updates
continuous
Settings...
setup
go
fitness
154000
0
0
231
ticks: 0
3D
generation
190
mean rho
-14.41168999
rho
1160
-1230
0
231
max-abs-prize
100
mutatn-rate
0.03
running mean rho
-19.7658144736842
Command Center
Clear
(turtle 17233): -2.0358447746279155E-123
(turtle 17233): -2.0083810806055562E-123
observer>

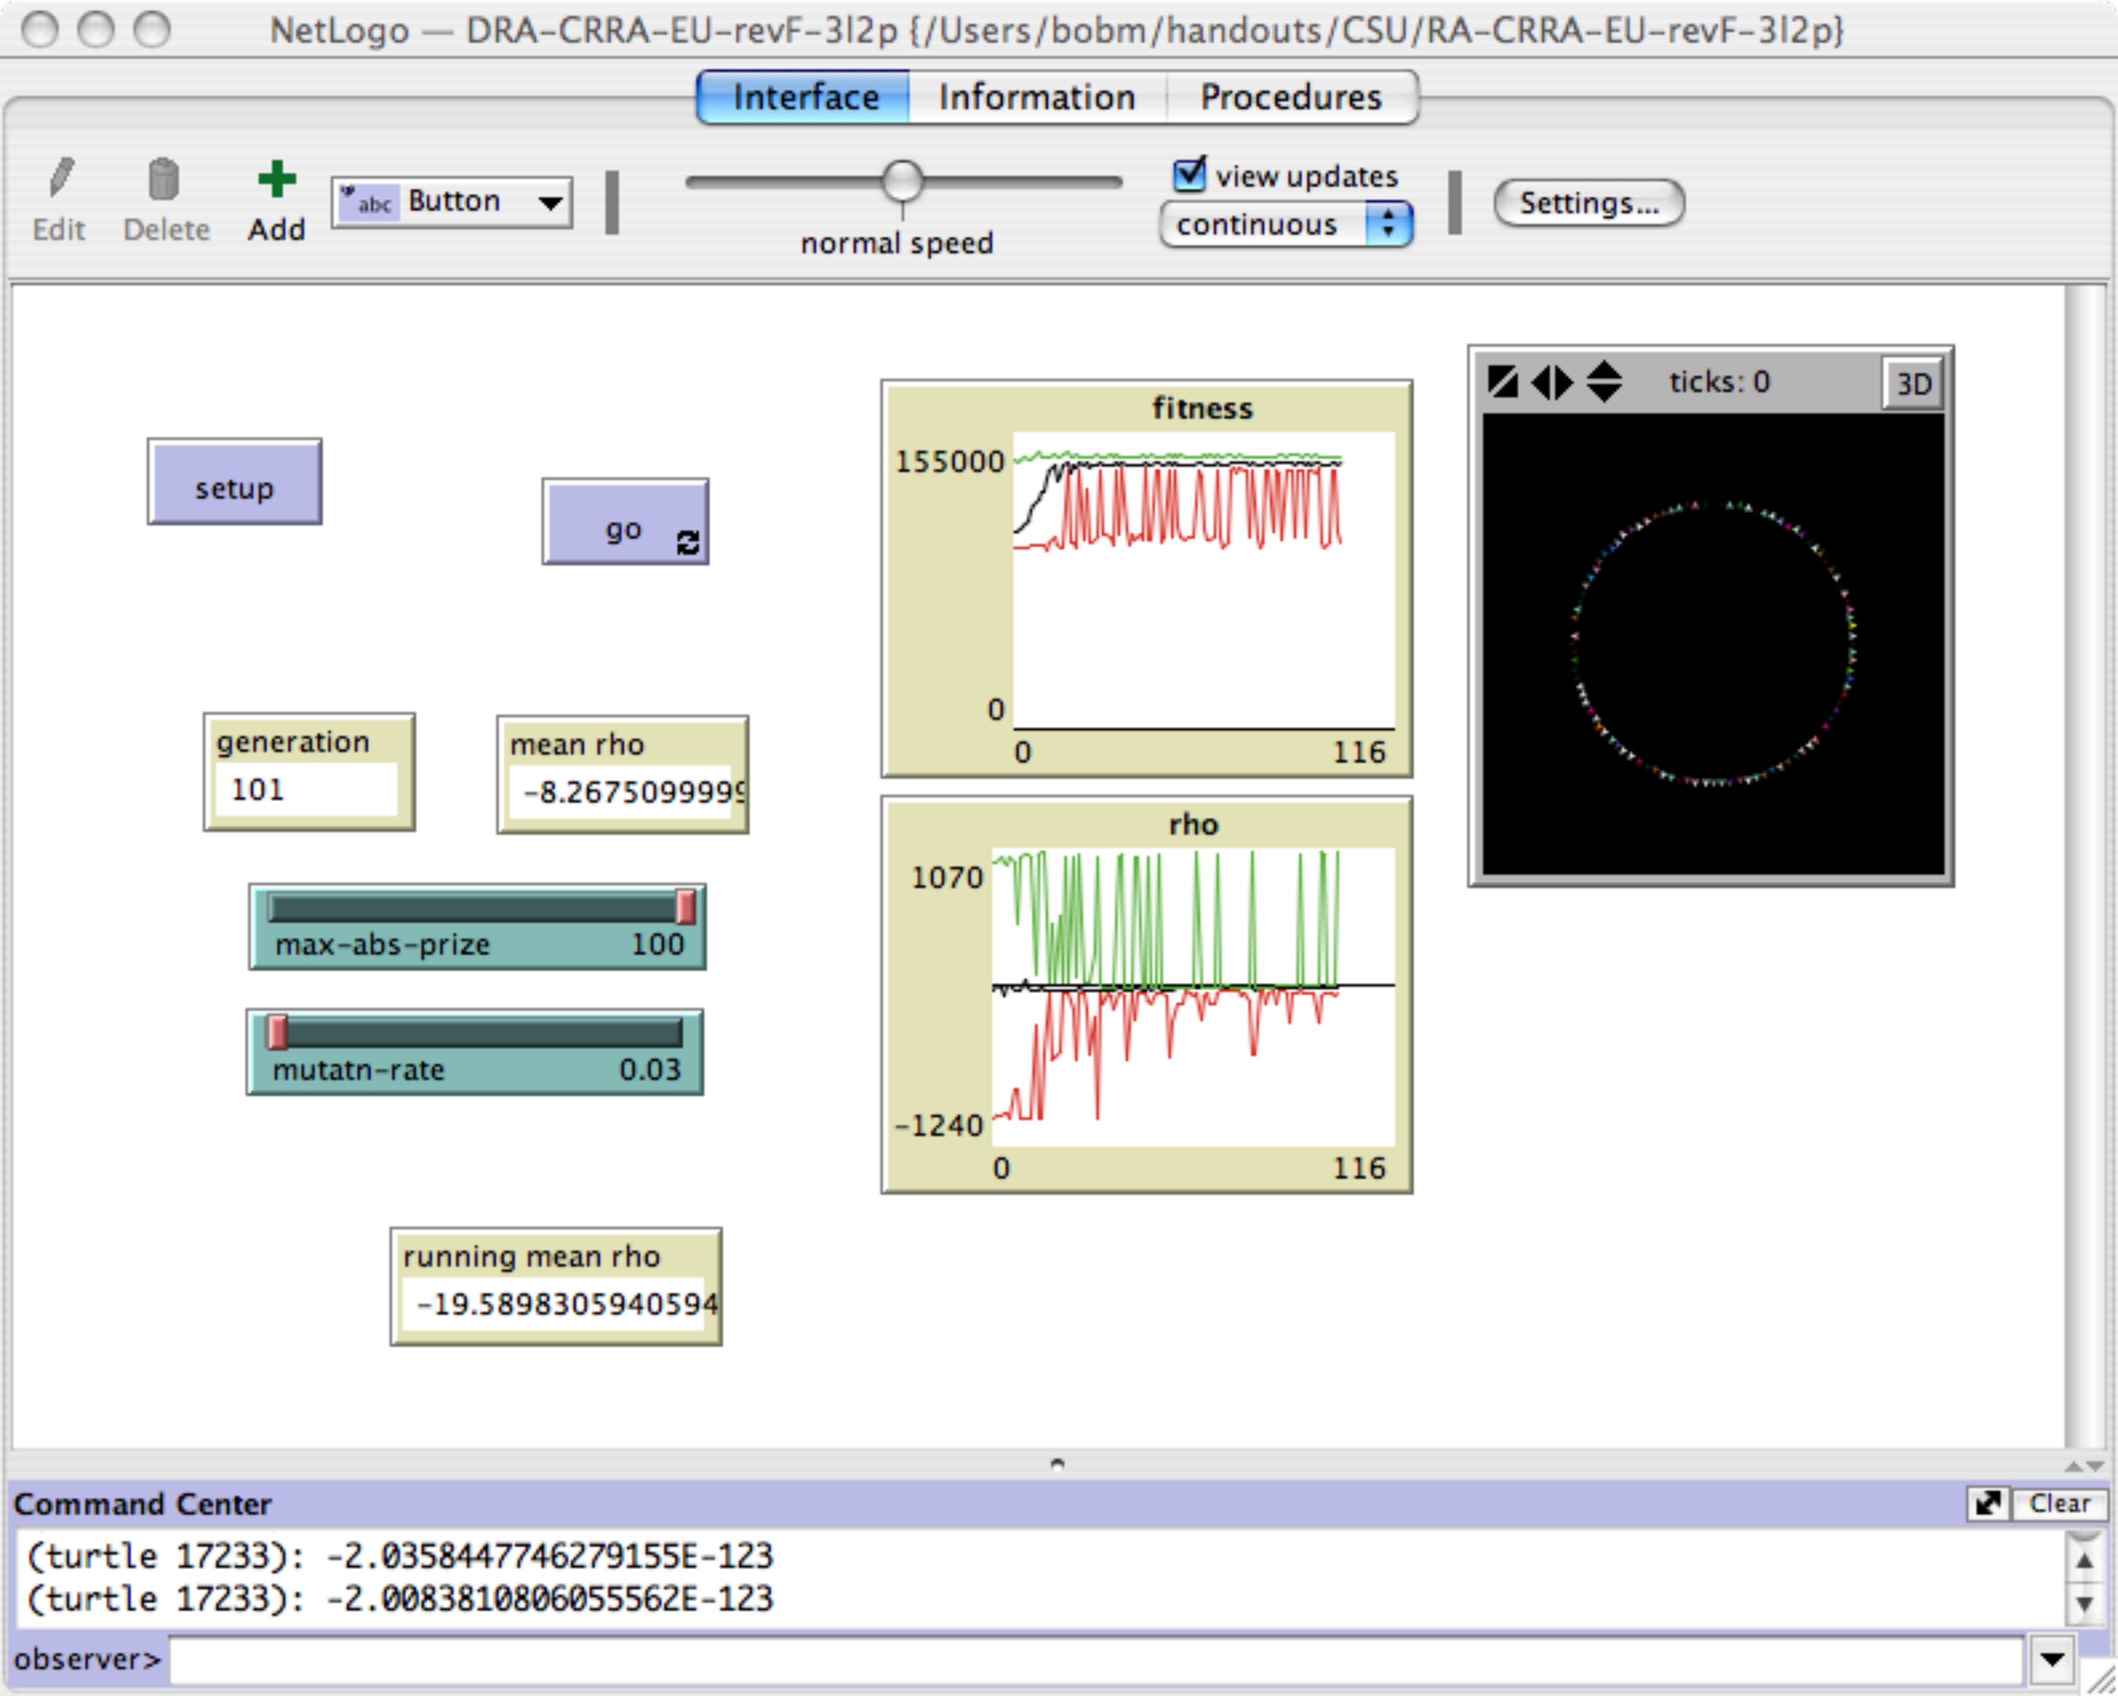
NetLogo — DRA-CRRA-EU-revF-3l2p {/Users/bobm/handouts/CSU/RA-CRRA-EU-revF-3l2p}
Interface
Information
Procedures
Edit
Delete
Add
Button
normal speed
view updates
continuous
Settings...
setup
go
fitness
155000
0
0
116
rho
1070
-1240
0
116
ticks: 0
3D
generation
101
mean rho
-8.2675099999
max-abs-prize
100
mutatn-rate
0.03
running mean rho
-19.5898305940594
Command Center
Clear
(turtle 17233): -2.0358447746279155E-123
(turtle 17233): -2.0083810806055562E-123
observer>

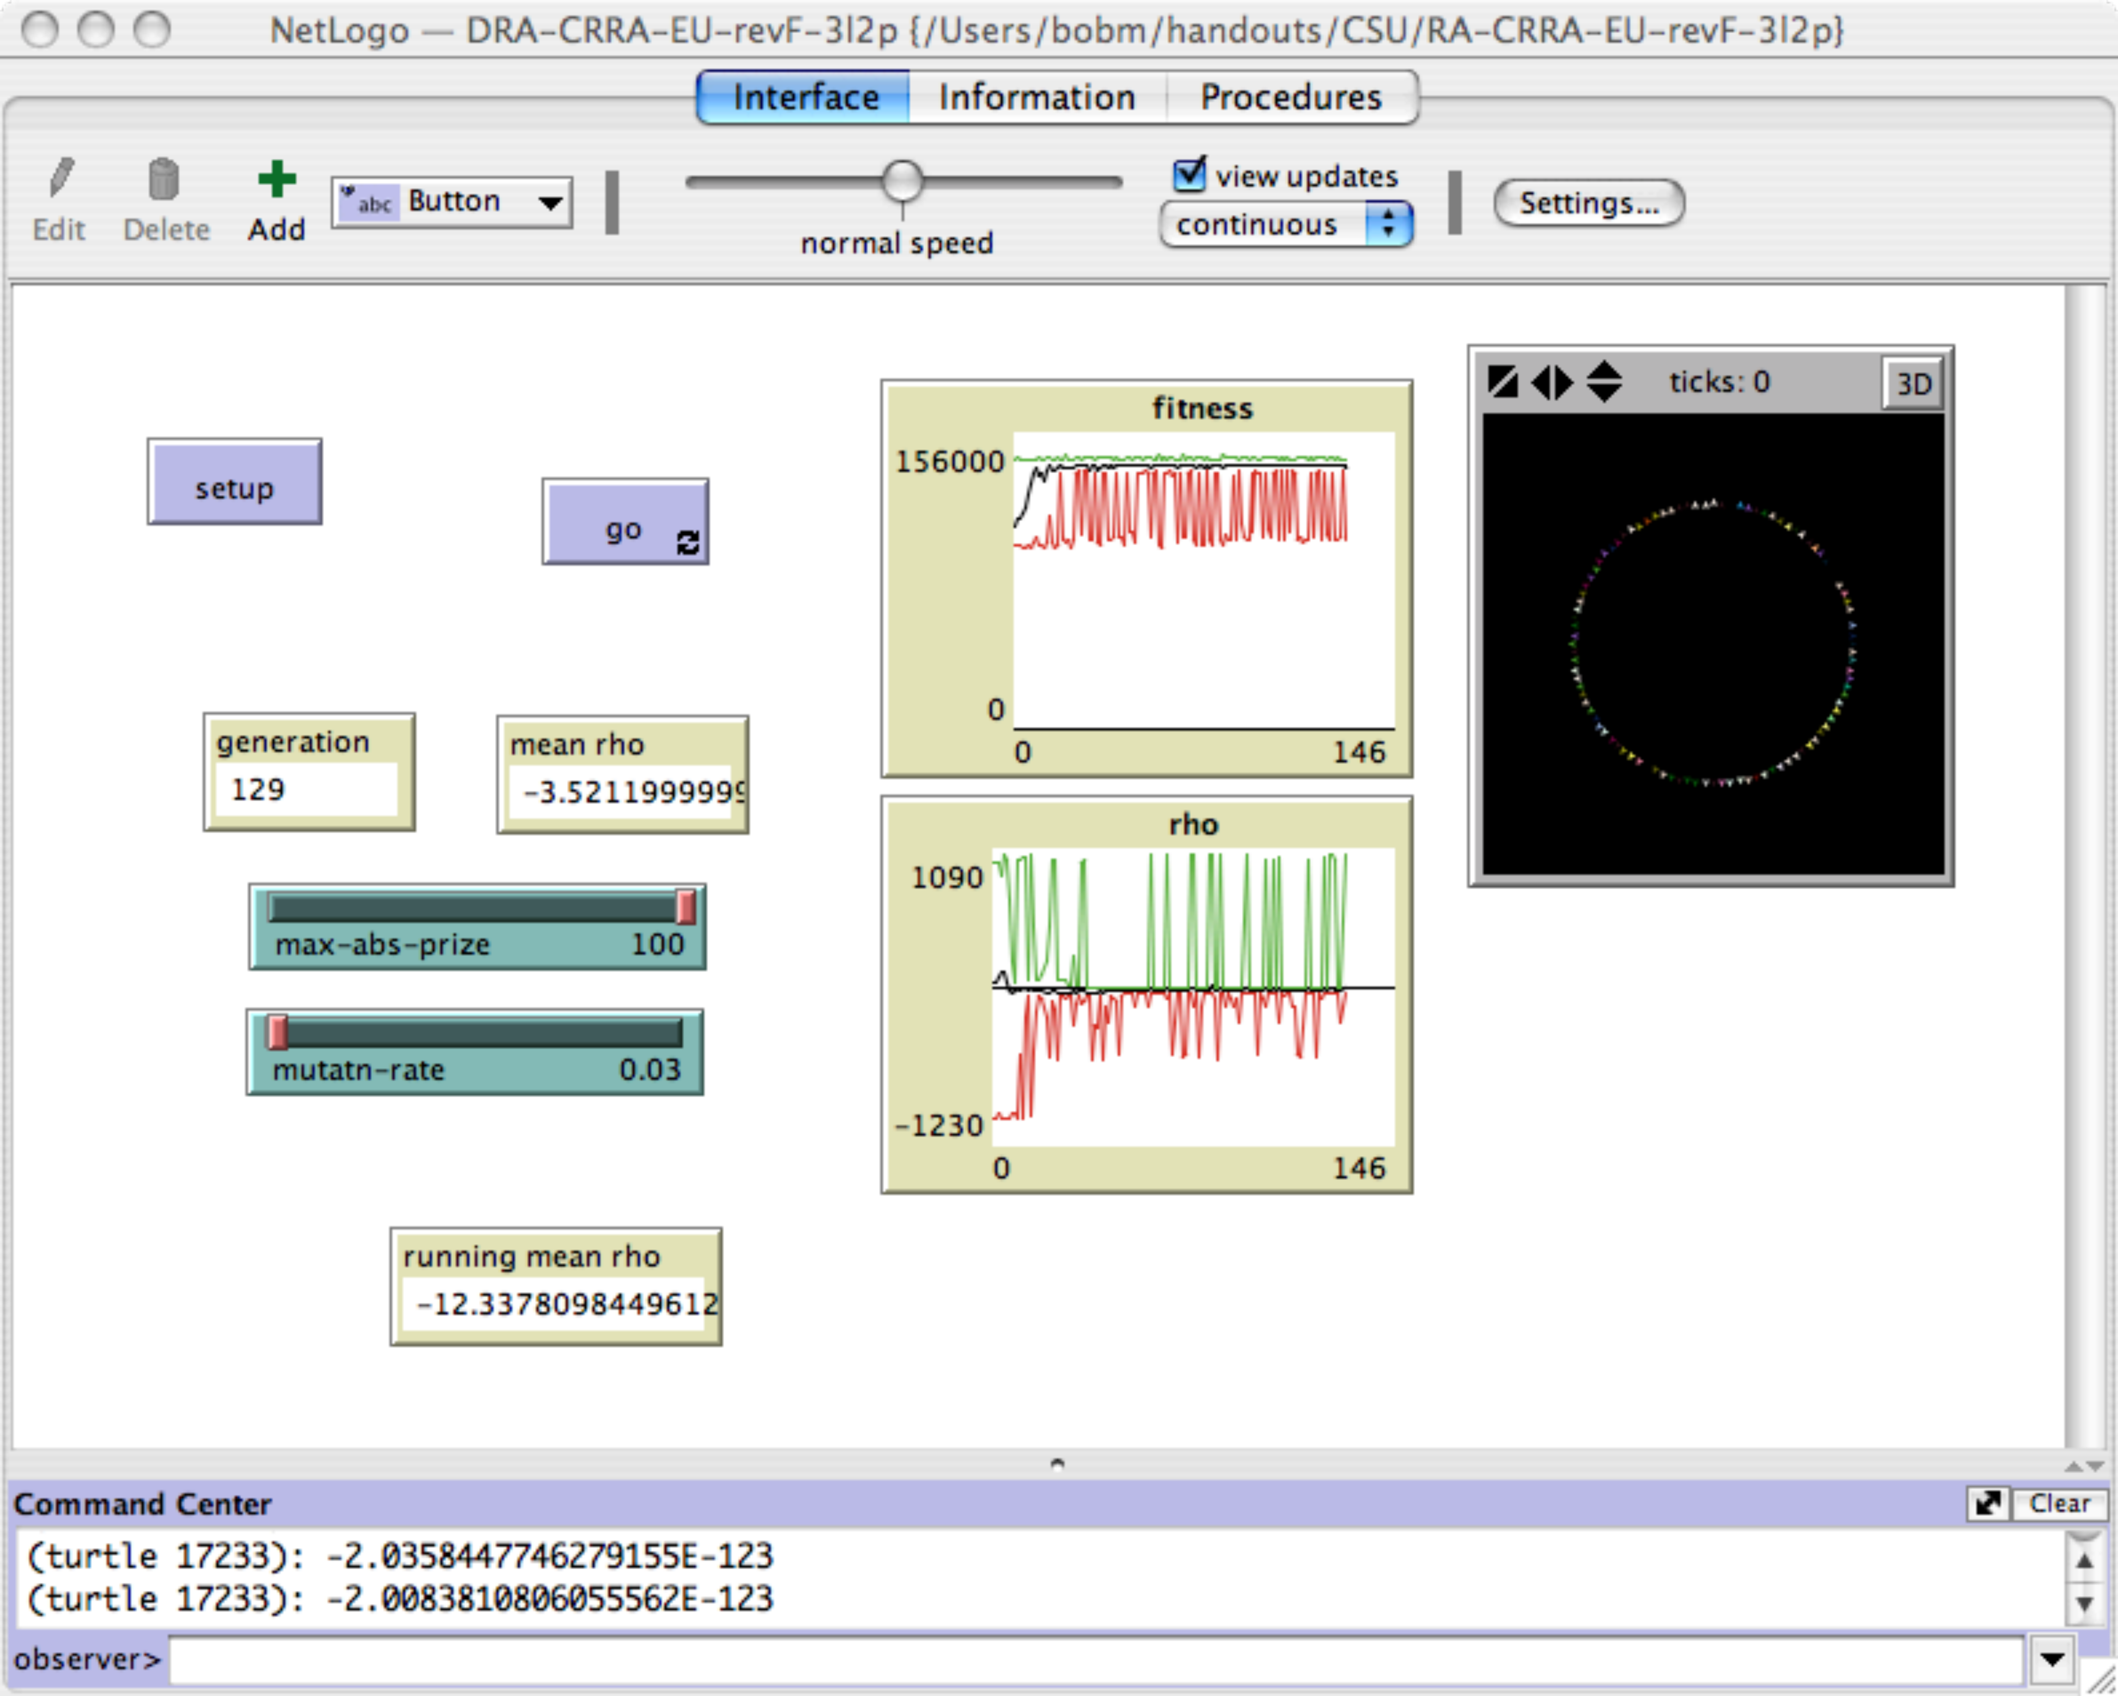
NetLogo — DRA-CRRA-EU-revF-3l2p {/Users/bobm/handouts/CSU/RA-CRRA-EU-revF-3l2p}
Interface
Information
Procedures
Edit
Delete
Add
Button
view updates
continuous
normal speed
Settings...
setup
go
fitness
156000
0
0
146
rho
1090
-1230
0
146
ticks: 0
3D
generation
129
mean rho
-3.521199999
max-abs-prize
100
mutatn-rate
0.03
running mean rho
-12.3378098449612
Command Center
Clear
(turtle 17233): -2.0358447746279155E-123
(turtle 17233): -2.0083810806055562E-123
observer>

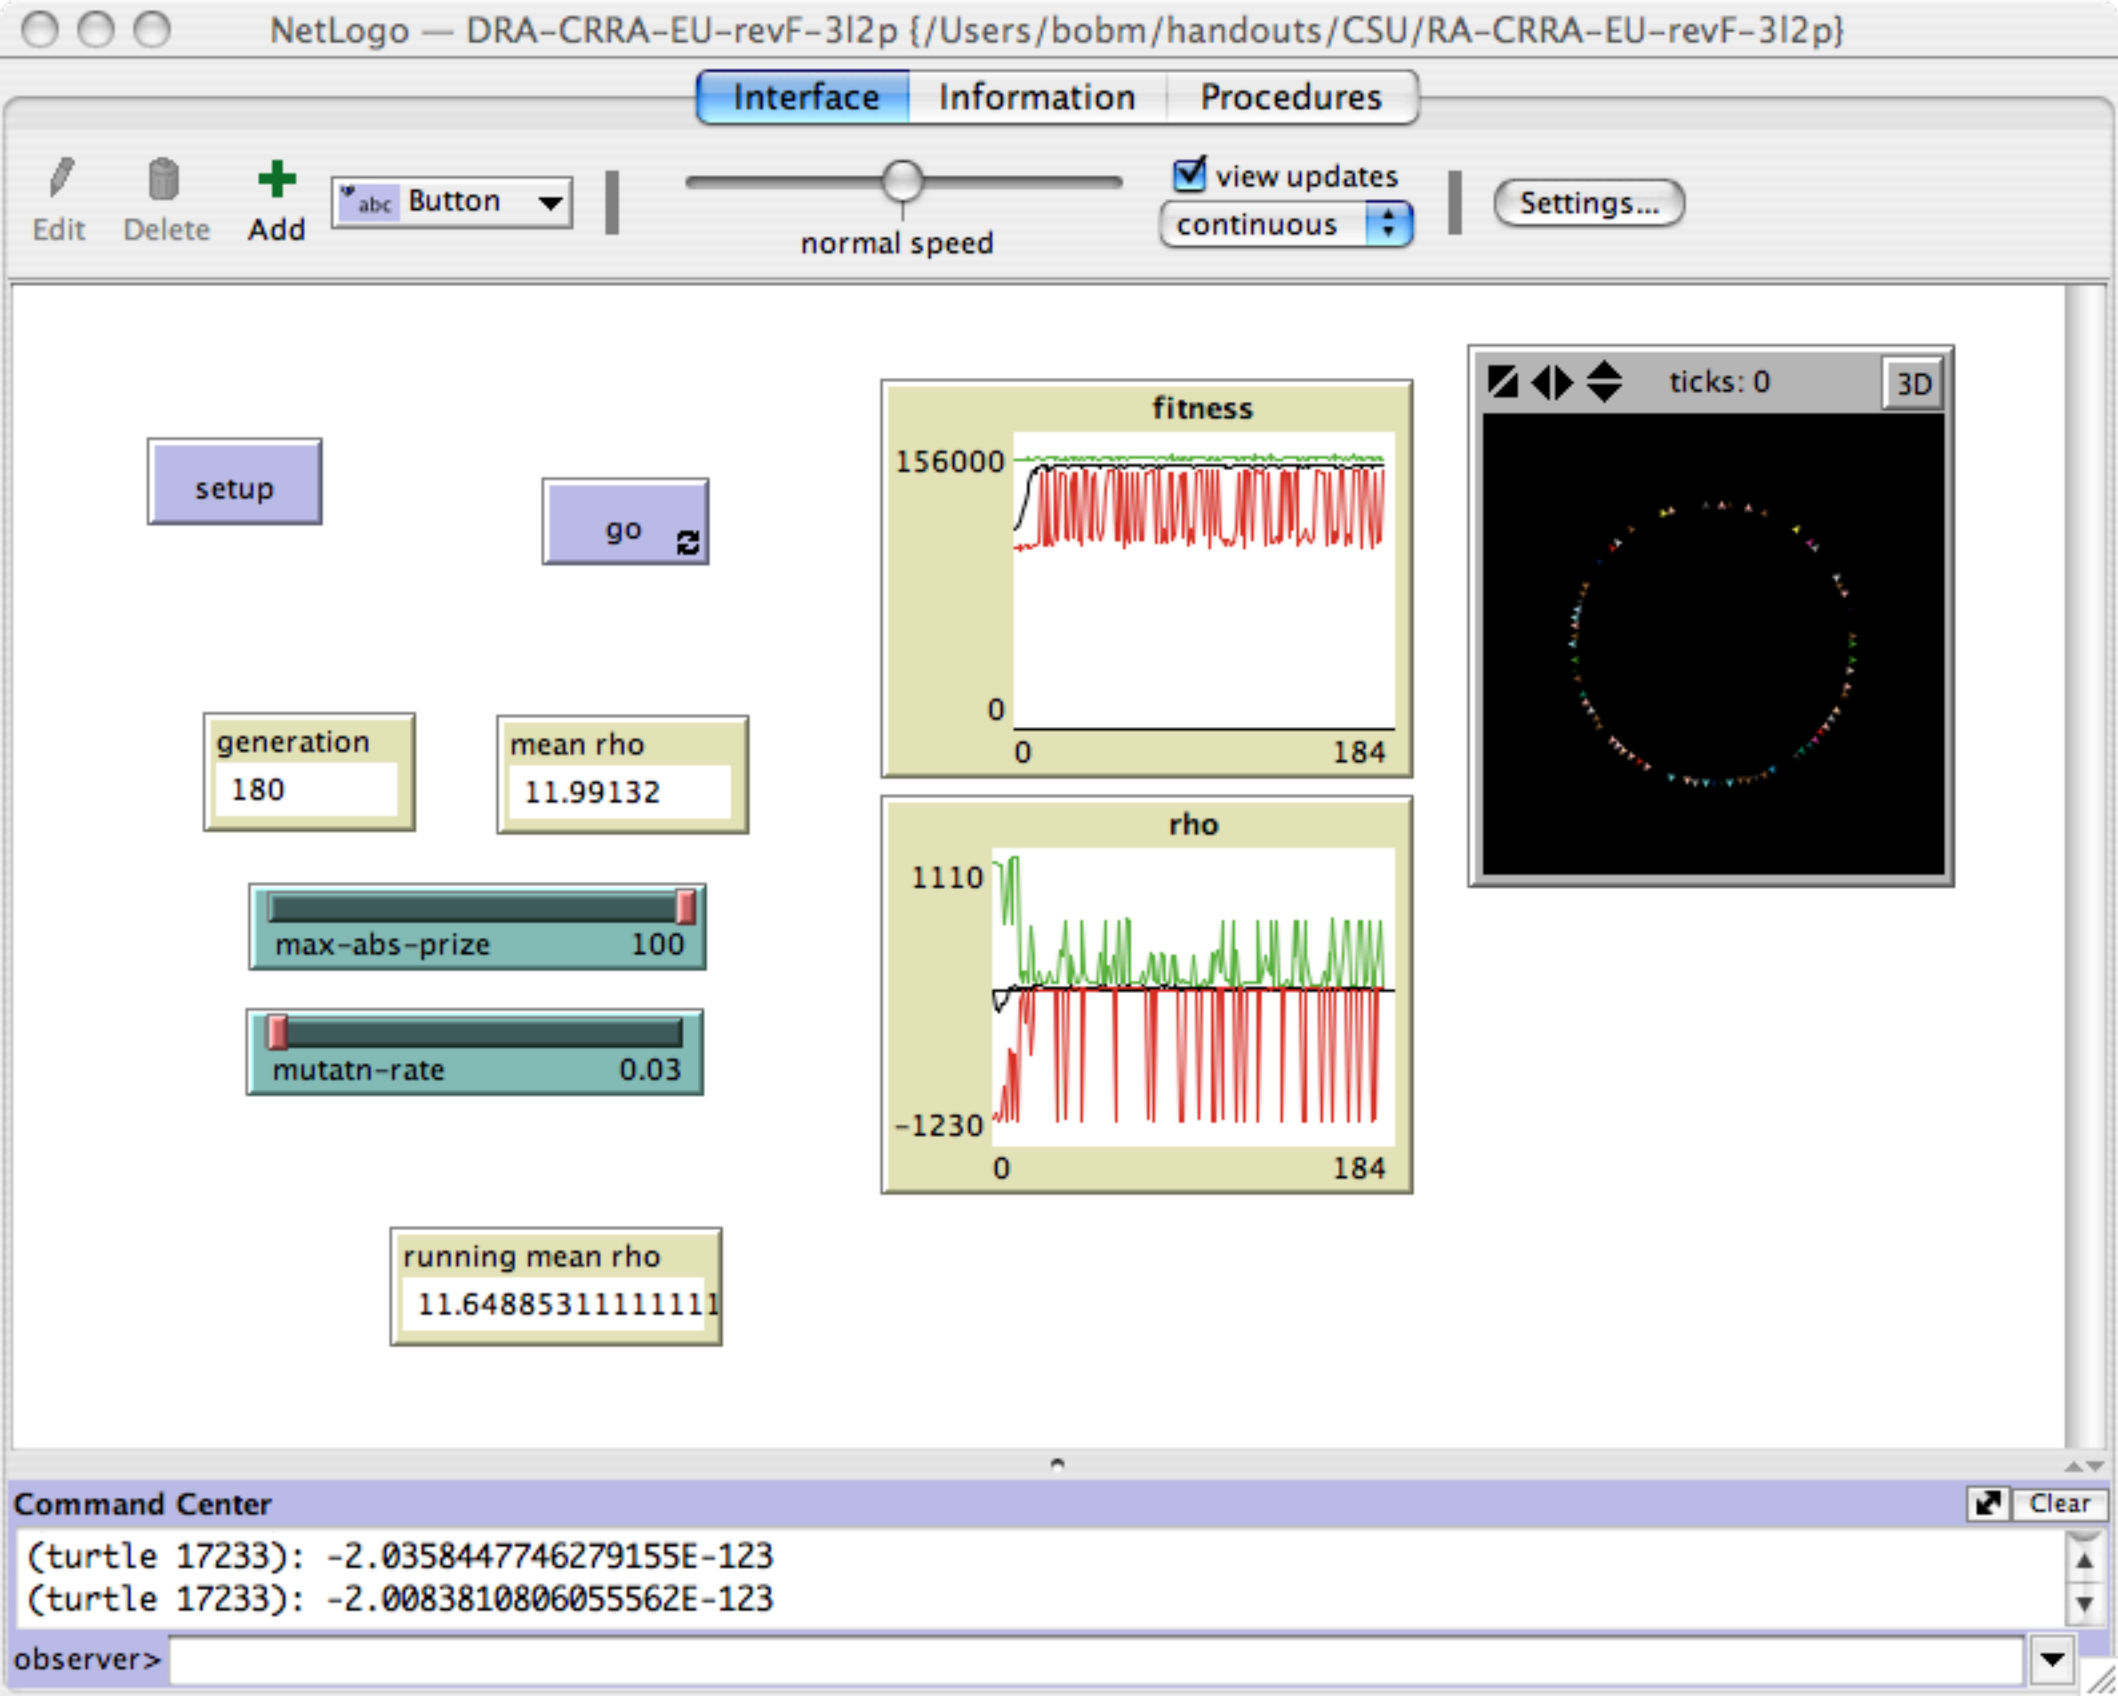
NetLogo — DRA-CRRA-EU-revF-3l2p {/Users/bobm/handouts/CSU/RA-CRRA-EU-revF-3l2p}
Interface
Information
Procedures
Edit
Delete
Add
Button
normal speed
view updates
continuous
Settings...
setup
go
generation
180
mean rho
11.99132
max-abs-prize
100
mutatn-rate
0.03
running mean rho
11.64885311111111
fitness
156000
0
0
184
rho
1110
-1230
0
184
ticks: 0
3D
Command Center
Clear
(turtle 17233): -2.0358447746279155E-123
(turtle 17233): -2.0083810806055562E-123
observer>

5. Discussion

Unlike the GA simulations of Szpiro (1997), we find that the best-performing CARA agents are risk-neutral, not risk averse. Because of the indirect way in which Szpiro modelled the risk profiles of his agents (unlike a referee's suggestion, footnote 3, Szpiro's model "only distinguishes between risk-averse automata and all others"), explanation of the contradictory results is not easy, but since our models allow any risk profile to emerge, we argue that they are more general than Szpiro's.

5. Discussion

Unlike the GA simulations of Szpiro (1997), we find that the best-performing CARA agents are risk-neutral, not risk averse. Because of the indirect way in which Szpiro modelled the risk profiles of his agents (unlike a referee's suggestion, footnote 3, Szpiro's model "only distinguishes between risk-averse automata and all others"), explanation of the contradictory results is not easy, but since our models allow any risk profile to emerge, we argue that they are more general than Szpiro's.

Should we be surprised that risk neutrality does better than risk aversion in CARA utility functions? Rabin (2000) suggests a reason why not, at least for small-stakes lotteries. He argues that von Neumann-Morgenstern expected-utility theory is inappropriate for reconciling actual human behaviour as revealed in risk attitudes over large stakes and small stakes. If there is risk aversion for small stakes, then expected-utility theory predicts wildly unrealistic risk aversion when the decision maker is faced with large stakes. Or risk aversion for large stakes must be accompanied by virtual risk neutrality for small stakes.

Rabin argues that *loss aversion* (Kahneman and Tversky 1979), rather than risk aversion, is a better (i.e. more realistic) explanation of how people actually behave when faced with risky decisions. This suggests possibilities for further simulations, although "loss aversion" suggests a prior conclusion.

Rabin argues that *loss aversion* (Kahneman and Tversky 1979), rather than risk aversion, is a better (i.e. more realistic) explanation of how people actually behave when faced with risky decisions. This suggests possibilities for further simulations, although “loss aversion” suggests a prior conclusion.

But we do not appeal to empirical evidence or even to prior beliefs of what sort of risk profile is best. Whereas there has been much research into reconciling actual human decision making with theory (see Arthur 1991), we are interested in seeing what is the best (i.e. most profitable) risk profile for agents faced with risky choices.

Rabin argues that *loss aversion* (Kahneman and Tversky 1979), rather than risk aversion, is a better (i.e. more realistic) explanation of how people actually behave when faced with risky decisions. This suggests possibilities for further simulations, although “loss aversion” suggests a prior conclusion.

But we do not appeal to empirical evidence or even to prior beliefs of what sort of risk profile is best. Whereas there has been much research into reconciling actual human decision making with theory (see Arthur 1991), we are interested in seeing what is the best (i.e. most profitable) risk profile for agents faced with risky choices.

And we find that for wealth-independent CARA utility functions (exponential) agents learn to become risk-neutral decision makers in order to maximise their returns when choosing among risky propositions. This is different from the risk-averse agents that Szpiro (1997) observed. But for wealth-dependent CRRA utility functions (CES) our agents often do learn to be slightly risk averse, as expected, but not always.

6. Bibliography

[1] Arrow, K., (1971) *Essays in the Theory of Risk Bearing*, Chicago: Markham.

[2] Arthur, W.B., (1991) "Designing economic agents that act like human agents: A behavioral approach to bounded rationality," *American Economic Review Papers & Proceedings*, 353–360.

[3] Gilbert, N. (2004) Axelrod's Iterated prisoners' dilemma tournament, from Chapter 10 of *Simulation for the Social Scientist*, by N. Gilbert and K.G. Troitzsch, 2nd ed., Open University Press, 2005.
http://cress.soc.surrey.ac.uk/s4ss/code/NetLogo/axelrod-ipd-ga.html

[4] Holland, J.H., (1992) *Adaptation in Natural and Artificial Systems*, 2nd Ed., Camb: MIT Press.

[5] Howard, R.A., (1968) "The foundations of decision analysis," *IEEE Trans. on Systems Science and Cybernetics*, ssc-4(3): 211–219.

[6] Kahneman, D. and A. Tversky, (1979) "Prospect theory: an analysis of decision under risk," *Econometrica*, 47, 263–291.

[7] Pratt, J.W., (1964) "Risk aversion in the small and in the large," *Econometerica,* 32: 122−136.

[8] Rabin, M., (2000) "Risk aversion and expected-utility theory: a calibration theorem," *Econometrica,* 68(5): 1281−1292.

[9] Szpiro, G.G., (1997) "The emergence of risk aversion," *Complexity*, 2(4): 31−39.

[10] von Neumann, J., and O. Morgenstern, (1944) *Theory of Games and Economic Behavior*, Princeton: Princeton University Press.

[11] Wilensky, U. (1999). NetLogo. http://ccl.northwestern.edu/netlogo. Center for Connected Learning and Computer-Based Modeling. Northwestern University, Evanston, IL.

Acknowlegements

I should like to thank Luis Izquierdo and the participants of the Complex Systems Research Summer School 2007 at Charles Sturt University, Jasmina Arifovic, Seth Tisue, and Nigel Gilbert for his implementation of the Genetic Algorithm in NetLogo.