

1 THE BUSINESS INVESTOR MODEL

This model was produced by S. Railsback and V. Grimm for Chapter 10 of the book *Agent-Based and Individual-Based Modeling*, 2nd edition (2019).

This is the first version of the model, as described in Section 10.4.1.

1.1 OVERVIEW

1.1.1 PURPOSE AND PATTERNS

The primary purpose of this model is to explore the effects of sensing—what information agents have and how they obtain it—on emergent outcomes of a model in which agents make adaptive decisions using sensed information. The model uses investment decisions as an example, but it is not intended to represent any real investment approach or business sector. (In fact, you will see by the end of chapter 12 that models like this one that are designed mainly to explore the system-level effects of sensing and other concepts can produce very different results depending on exactly how those concepts are implemented. What can we learn from such models if their results depend on what seem like details? In part III of this course we will learn to solve this problem by tying models more closely to real systems.)

This model can be thought of as approximately representing people who buy and operate local businesses: it assumes investors are familiar with businesses they could buy within a limited range of their own experience. The investment “environment” assumes that there is no cost of entering or switching businesses (e.g., as if capital to buy a business is borrowed and the repayment is included in the annual profit calculation), that high profits are rarer than low profits, and that risk of failure is unrelated to profit. We can use the model to address questions such as how the average wealth of the investors, and how evenly wealth is distributed among individuals, depends on how much information the investors can sense.

Because this model is conceptual and not intended to represent a specific real system, we can only use general patterns as criteria for its usefulness. These patterns are that investor decisions depend on the profit and risk “landscape,” and that the decisions change with investor wealth.

1.1.2 ENTITIES, STATE VARIABLES, AND SCALES

The entities in this model are investor agents (turtles) and business alternatives (patches) that vary in profit and risk. The investors have state variables for their location in the space and for their current wealth (W , in money units).

The landscape is a grid of business patches, each of which has two static variables: the annual net profit that a business there would provide (P , in money units such as dollars per year) and the annual risk of that business failing and its investor losing all its wealth (F , as probability per year). This landscape is 19×19 patches in size with no wrapping at its edges.

The model time step is 1 year, and simulations run for 50 years.

1.1.3 PROCESS OVERVIEW AND SCHEDULING

The model includes the following actions that are executed in this order each time step:

Investor repositioning. The investors decide whether any similar business (adjacent patch) offers a better tradeoff of profit and risk; if so, they “reposition” and transfer their investment to that patch, by moving there. The repositioning submodel is described at element 7. Only one investor can occupy a patch at a time. The agents execute this repositioning action in randomized order.

Accounting. The investors update their wealth state variable. W is set equal to the previous wealth plus the profit of the agent’s current patch. However, unexpected failure of the business is also included in the accounting action. This event is a stochastic function of F at the investor’s patch. If a uniform random number between zero and one is less than F , then the business fails: the investor’s W is set to zero, but the investor stays in the model and continues to behave in the same way.

Output. The View, plots, and an output file are updated.

1.2 DESIGN CONCEPTS

Basic principles. The basic topic of this model is how agents make decisions involving tradeoffs between several objectives—here, increasing profit and decreasing risk.

Emergence. The model’s primary output is mean investor wealth over time. Important secondary outputs are the mean profit and risk chosen by investors over time, and the number of investors who have suffered a failure. These outputs emerge from how individual investors make their tradeoff between profit and risk but also from the “business climate”: the ranges of P and F values among patches and the number of investors competing for locations on the landscape.

Adaptive behavior. The adaptive behavior of investor agents is repositioning: the decision of which neighboring business to move to (or whether to stay put), considering the profit and risk of these alternatives. Each time step, investors can reposition to any unoccupied one of their adjacent patches or retain their current position. In this version of the model, investors use a simplified microeconomic analysis to make their decision, moving to the patch providing the highest value of an objective function.

Objective. (In economics, the term utility is used for the objective that agents seek.) Investors rate business alternatives by a utility measure that represents their expected future wealth at the end of a time horizon (T , a number of future years; we use 5) if they buy and operate the business. This expected future wealth is a function of the investor’s current wealth and the profit and failure risk offered by the patch:

$$U = (W + TP)(1 - F)^T$$

where U is the expected utility for the patch, W is the investor’s current wealth, and P and F are defined above. The term $(W + TP)$ estimates investor wealth at the end of the time horizon if no failures occur. The term $(1 - F)^T$ is the probability of not having a failure over the time horizon; it reduces utility more as failure risk increases. (Economists might expect to use a utility measure such as present value that includes a discount rate to reduce the value of future profit. We ignore discounting to keep this model simple.)

Prediction. The utility measure estimates utility over a time horizon by using the explicit prediction that P and F will remain constant over the time horizon. This assumption is accurate here because the patches’ P and F values are static.

Sensing. The investor agents are assumed to know the profit and risk at their own patch and the adjacent neighbor patches, without error.

Interaction. The investors interact with each other only indirectly via competition for patches; an investor cannot take over a business (move into a patch) that is already occupied by another investor. Investors execute their repositioning action in randomized order, so there is no hierarchy in this competition: investors with higher wealth have no advantage over others in competing for locations.

Stochasticity. The initial state of the model is stochastic: the values of P and F of each patch, and initial investor locations, are set randomly. *Stochasticity* is thus used to simulate an investment environment where alternatives are highly variable and risk is not correlated with profit. The values of P are drawn from an exponential distribution, which produces many patches with low profits and a few patches with high profits (we will learn more about this and other random number distributions in chapter 15). The random exponential distribution of P makes the results of this model especially variable, even among model runs with the same parameter values. Whether each investor fails each year is also stochastic, a simple way to represent risk. The investor reposition action uses stochasticity only in the very unlikely event that more than one potential destination patch offers the same highest utility; when there is such a tie, the agent randomly chooses one of the tied patches to move to.

Observation. The View shows the location of each agent on the investment landscape. Having the investors put their pen down lets us observe how many patches each has used. Graphs show the mean profit and risk experienced by investors, and mean investor wealth over time. The standard deviation in wealth among

investors is also graphed as a simple and appropriate measure of how evenly wealth is distributed among investors.

Learning and *collectives* are not represented.

1.3 DETAILS

1.3.1 INITIALIZATION

The value of P for each patch is drawn from a random exponential distribution with a mean of 5000. The value of F for each patch is drawn randomly from a uniform real number distribution with a minimum of 0.01 and a maximum of 0.1.

Twenty-five investor agents are initialized and put in random patches, but investors cannot be placed in a patch already occupied by another investor. Their wealth state variable W is initialized to zero.

1.3.2 INPUT DATA

No time-series inputs are used.

1.3.3 SUBMODELS

Investor repositioning. An investor identifies all the businesses that it could invest in: any of the neighboring eight (or fewer if on the edge of the space) patches that are unoccupied, plus its current patch. The investor then determines which of these alternatives provides the highest value of the utility function, and moves (or stays) there.

Accounting. This action is fully described above (“Process overview and scheduling”).