The Neuroscience of Human Decision Making

Dr. Olav E. Krigolson

Associate Director

Centre for Biomedical Research

University of Victoria

Krigolson Lab

Theoretical

Applied

Neuroscience

Decision Making Theory

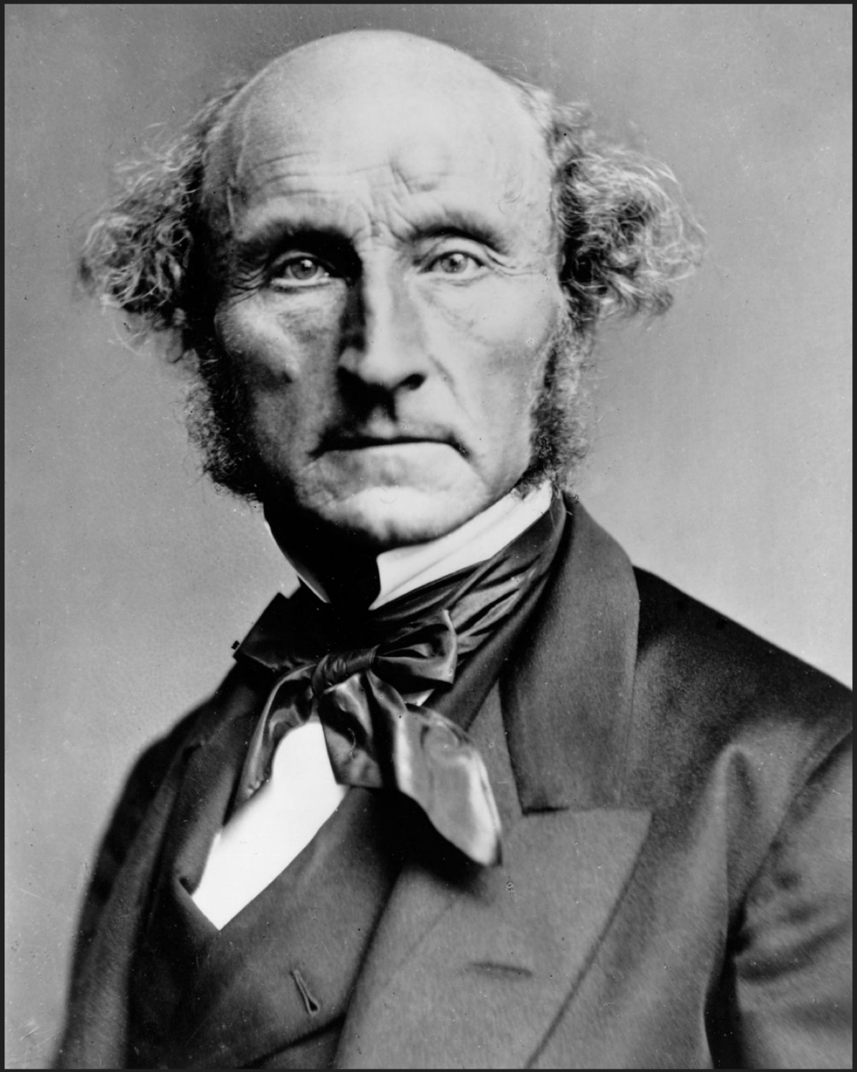

Utilitarianism

People seek actions
that increase utility
and avoid actions that
decrease utility

Mill, 1861

Decision Making

Our ability to process multiple alternatives and choose the option that maximizes utility

Huygens, 1657

Expected Value = Value x Probability

Expanded Form

$$EV = Gain \times P_G - Cost \times P_C$$

A Sample Problem

Problem 1

Would you play a gamble that has a 40% chance to win $1000 or a 70% chance to win $600?

Random Variable

A random variable *x* takes on a defined set of values with different probabilities.

- For example, if you roll a die, the outcome is random (not fixed) and there are 6 possible outcomes, each of which occur with probability one-sixth.
- For example, if you poll people about their voting preferences, the percentage of the sample that responds "Yes on Proposition 100" is a also a random variable (the percentage will be slightly differently every time you poll).

Roughly, <u>probability</u> is how frequently we expect different outcomes to occur if we repeat the experiment over and over.

Random variables can be discrete or continuous

- **Discrete** random variables have a countable number of outcomes
  - Examples: Dead/alive, treatment/placebo, dice, counts, etc.

- **Continuous** random variables have an infinite continuum of possible values.
  - Examples: blood pressure, weight, the speed of a car, the real numbers from 1 to 6.

Probability Functions

- A probability function maps the possible values of x against their respective probabilities of occurrence, $p(x)$

- $p(x)$ is a number from 0 to 1.0.

- The area under a probability function is always 1.

Discrete example: Roll of a Die

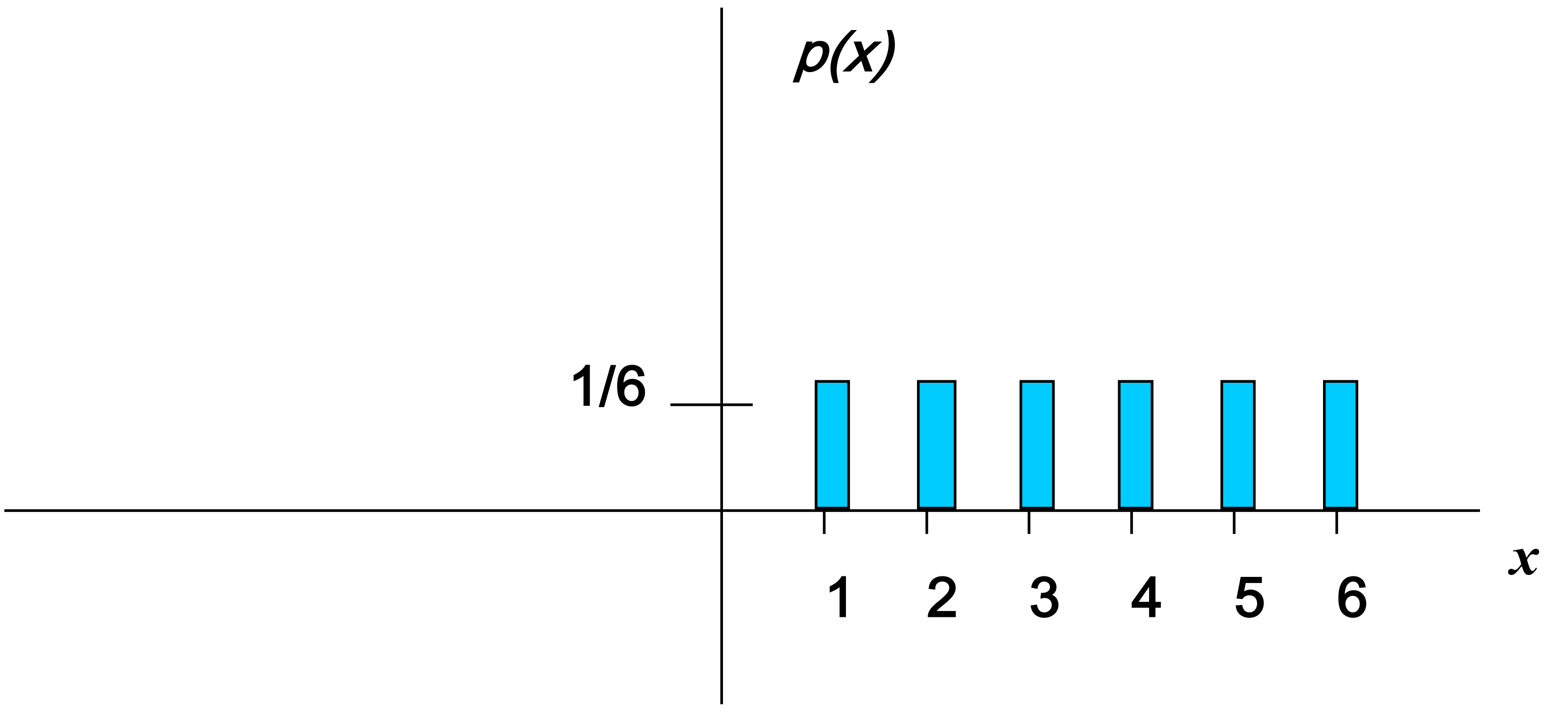

$$\sum_{\text{all } x} P(x) = 1$$

Cumulative Distribution Function (CDF)

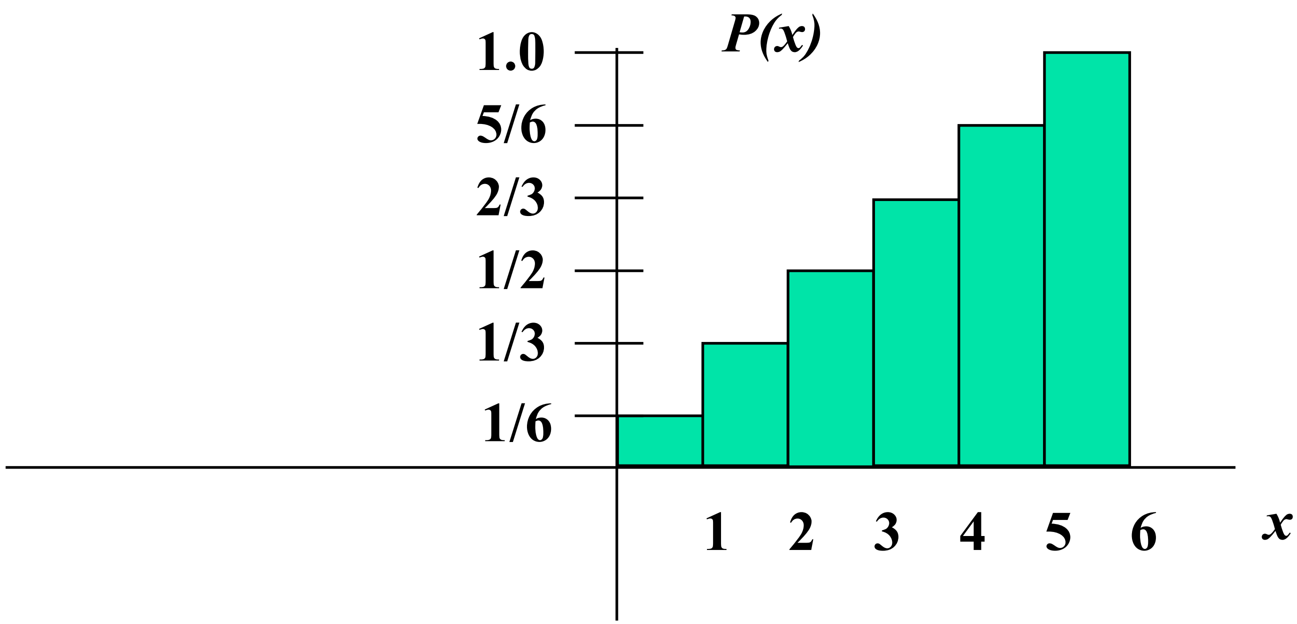

Practice Problem:

- The number of patients seen in the ER in any given hour is a random variable represented by *x*. The probability distribution for *x* is:

x	10	11	12	13	14
P(x)	.4	.2	.2	.1	.1

Find the probability that in a given hour:

a. exactly 14 patients arrive $p(x=14) = .1$

b. At least 12 patients arrive $p(x \geq 12) = (.2 + .1 + .1) = .4$

c. At most 11 patients arrive $p(x \leq 11) = (.4 + .2) = .6$

Review Question 1

If you toss a die, what's the probability that you roll a 3 or less?

a. 1/6
b. 1/3
c. 1/2
d. 5/6
e. 1.0

Review Question 1

If you toss a die, what’s the probability that you roll a 3 or less?

a. 1/6
b. 1/3
c. 1/2
d. 5/6
e. 1.0

Review Question 2

Two dice are rolled and the sum of the face values is six? What is the probability that at least one of the dice came up a 3?

a. 1/5
b. 2/3
c. 1/2
d. 5/6
e. 1.0

Review Question 2

Two dice are rolled and the sum of the face values is six. What is the probability that at least one of the dice came up a 3?

a. **1/5**
b. 2/3
c. 1/2
d. 5/6
e. 1.0

How can you get a 6 on two dice? 1-5, 5-1, 2-4, 4-2, 3-3

One of these five has a 3.

$\therefore$ **1/5**

For example, the probability of x falling within 1 to 2:

Clinical Example: Survival times after lung transplant may roughly follow an exponential function.

Then, the probability that a patient will die in the second year after surgery (between years 1 and 2) is 23%.

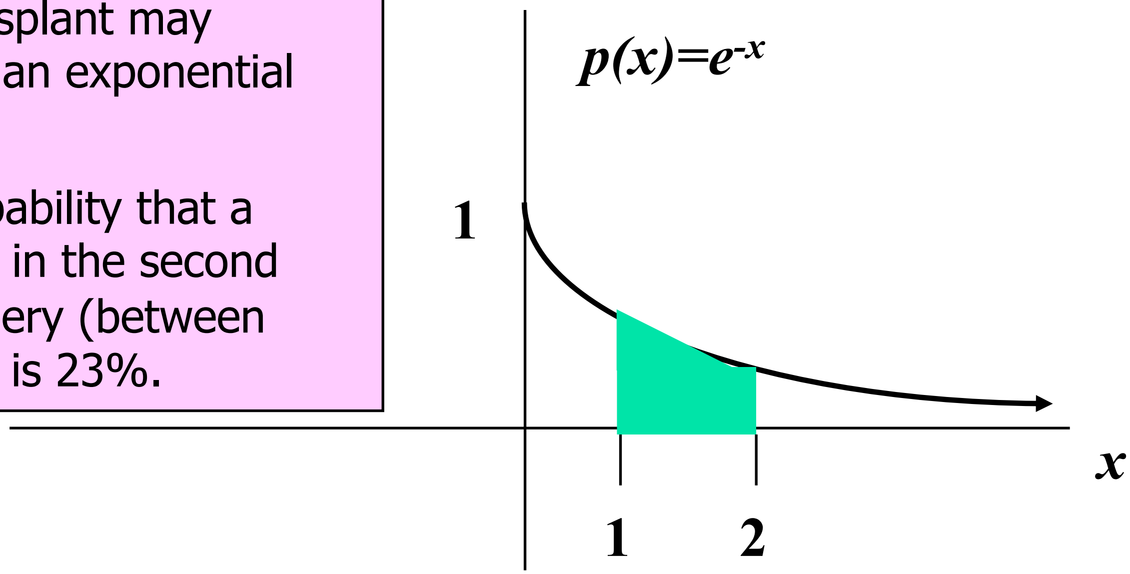

$$P(1 \leq x \leq 2) = \int_1^2 e^{-x} = -e^{-x} \Big|_1^2 = -e^{-2} - -e^{-1} = -.135 + .368 = .23$$

Back to Expected Value

Example: the Lottery

- <u>The Lottery</u> (also known as a tax on people who are bad at math…)
- A certain lottery works by picking 6 numbers from 1 to 49. It costs $1.00 to play the lottery, and if you win, you win $2 million after taxes.
- *If you play the lottery once, what are your expected winnings or losses?*

Lottery

Calculate the probability of winning in 1 try:

$$\frac{1}{\binom{49}{6}} = \frac{1}{\frac{49!}{43!6!}} = \frac{1}{13{,}983{,}816} = 7.2 \text{ x } 10^{-8}$$

"49 choose 6"

Out of 49 numbers, this is the number of distinct combinations of 6.

The probability function (note, sums to 1.0):

x$	*p(x)*
-1	.999999928
+ 2 million	*7.2 x* 10^{-8}

Expected Value

The probability function

x$	*p(x)*
-1	.999999928
+ 2 million	7.2×10^{-8}

Expected Value

E(X) = P(win)*$2,000,000 + P(lose)*-$1.00

$= 2.0 \times 10^{6} * 7.2 \times 10^{-8} + .999999928\,(-1) = .144 - .999999928 = -\$.86$

Negative expected value is never good!

You shouldn't play if you expect to lose money!

Expected Value

If you play the lottery every week for 10 years, what are your expected winnings or losses?

520 x (-.86) = -$447.20

Why casinos give out free drinks

A roulette wheel has the numbers 1 through 36, as well as 0 and 00. If you bet \$1 that an odd number comes up, you win or lose \$1 according to whether or not that event occurs. If random variable X denotes your net gain, $X=1$ with probability 18/38 and $X=-1$ with probability 20/38.

$E(X) = 1(18/38) - 1\,(20/38) = -\$.053$

On average, the casino wins (and the player loses) 5 cents per game.

The casino rakes in even more if the stakes are higher:

$E(X) = 10(18/38) - 10\,(20/38) = -\$.53$

If the cost is \$10 per game, the casino wins an average of 53 cents per game. If 10,000 games are played in a night, that's a cool \$5300.

Another example of Expected Value

- Take the show “Deal or No Deal”
- Everyone know the rules?
- Let’s say you are down to two cases left. $1 and $400,000. The banker offers you $225,000.
- So, Deal or No Deal?

Deal or No Deal…

x$	*p(x)*
+1	.50
+$400,000	.50

x$	*p(x)*
+$225,000	1.0

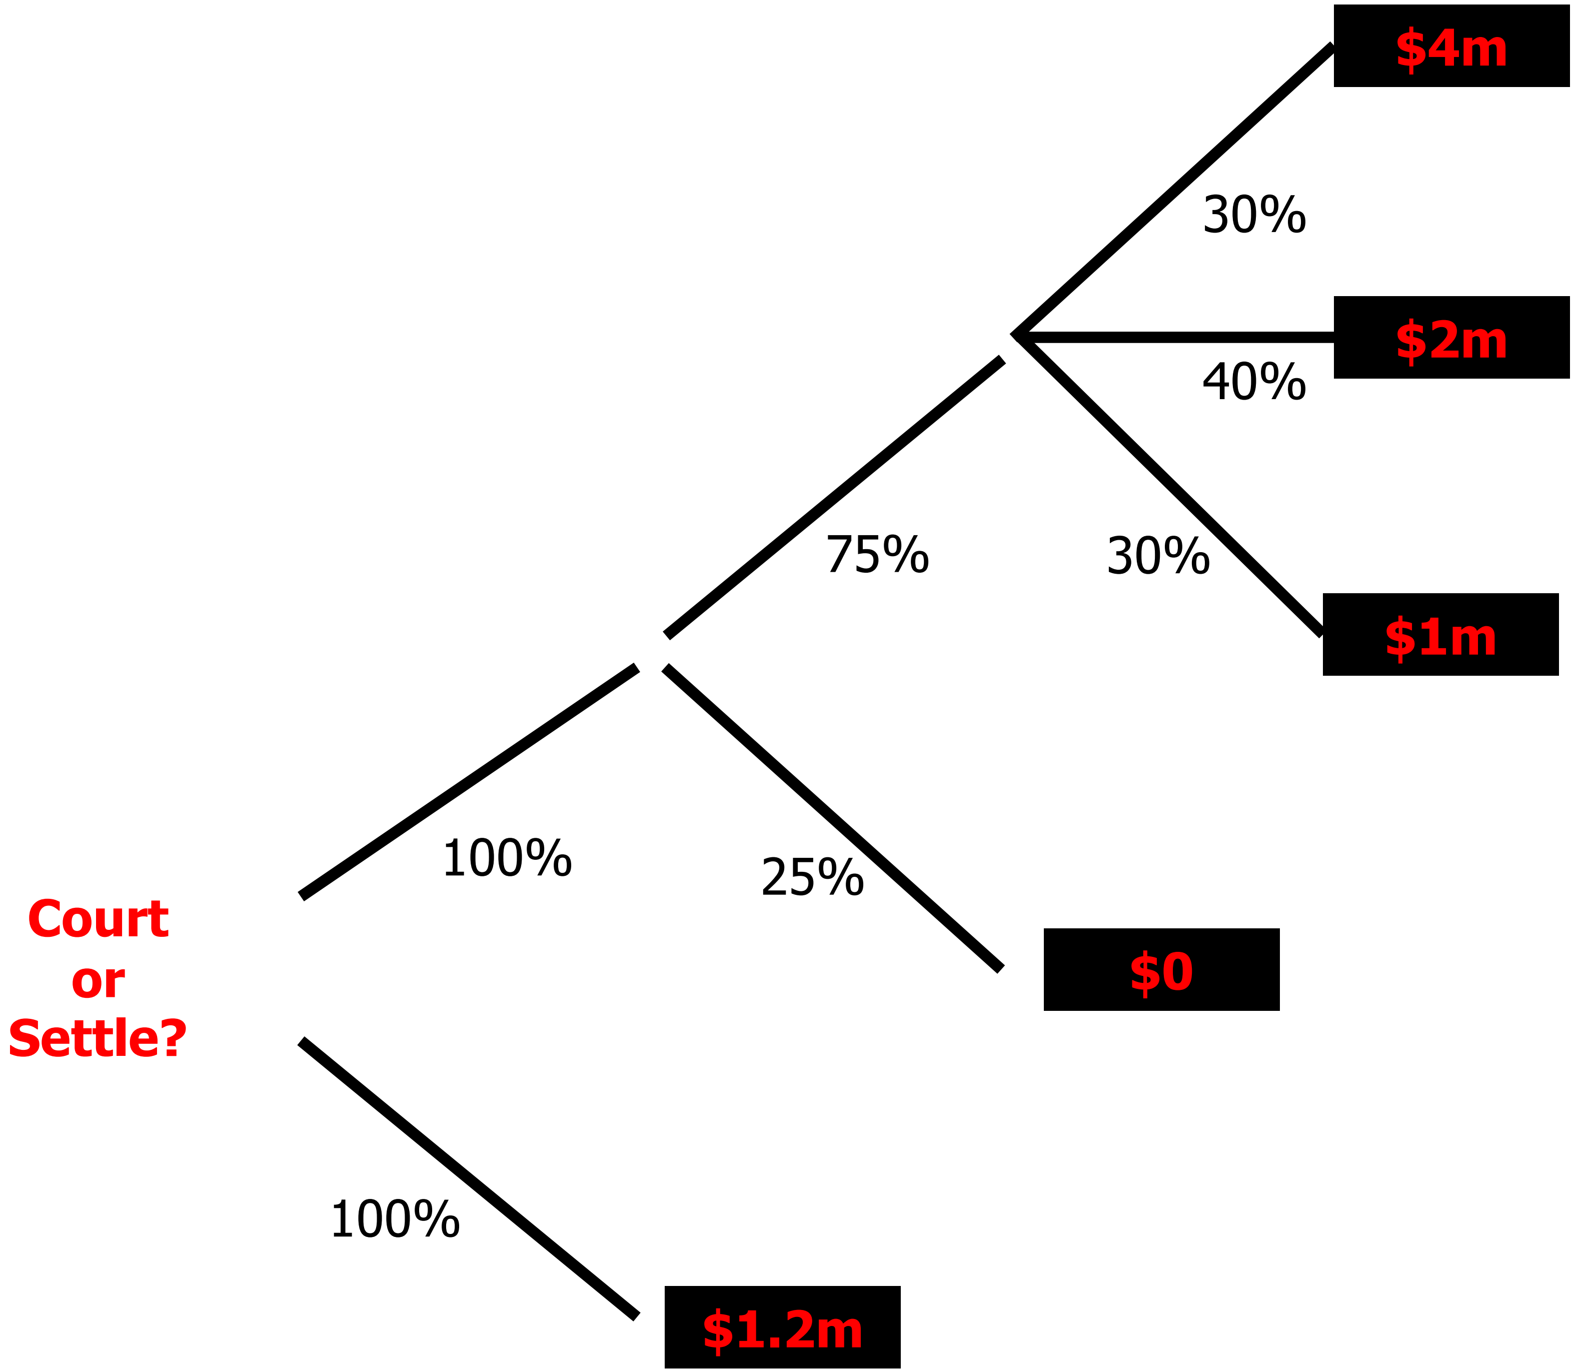
$4m
30%
$2m
40%
75%
30%
$1m
100%
25%
$0
Court
or
Settle?
100%
$1.2m

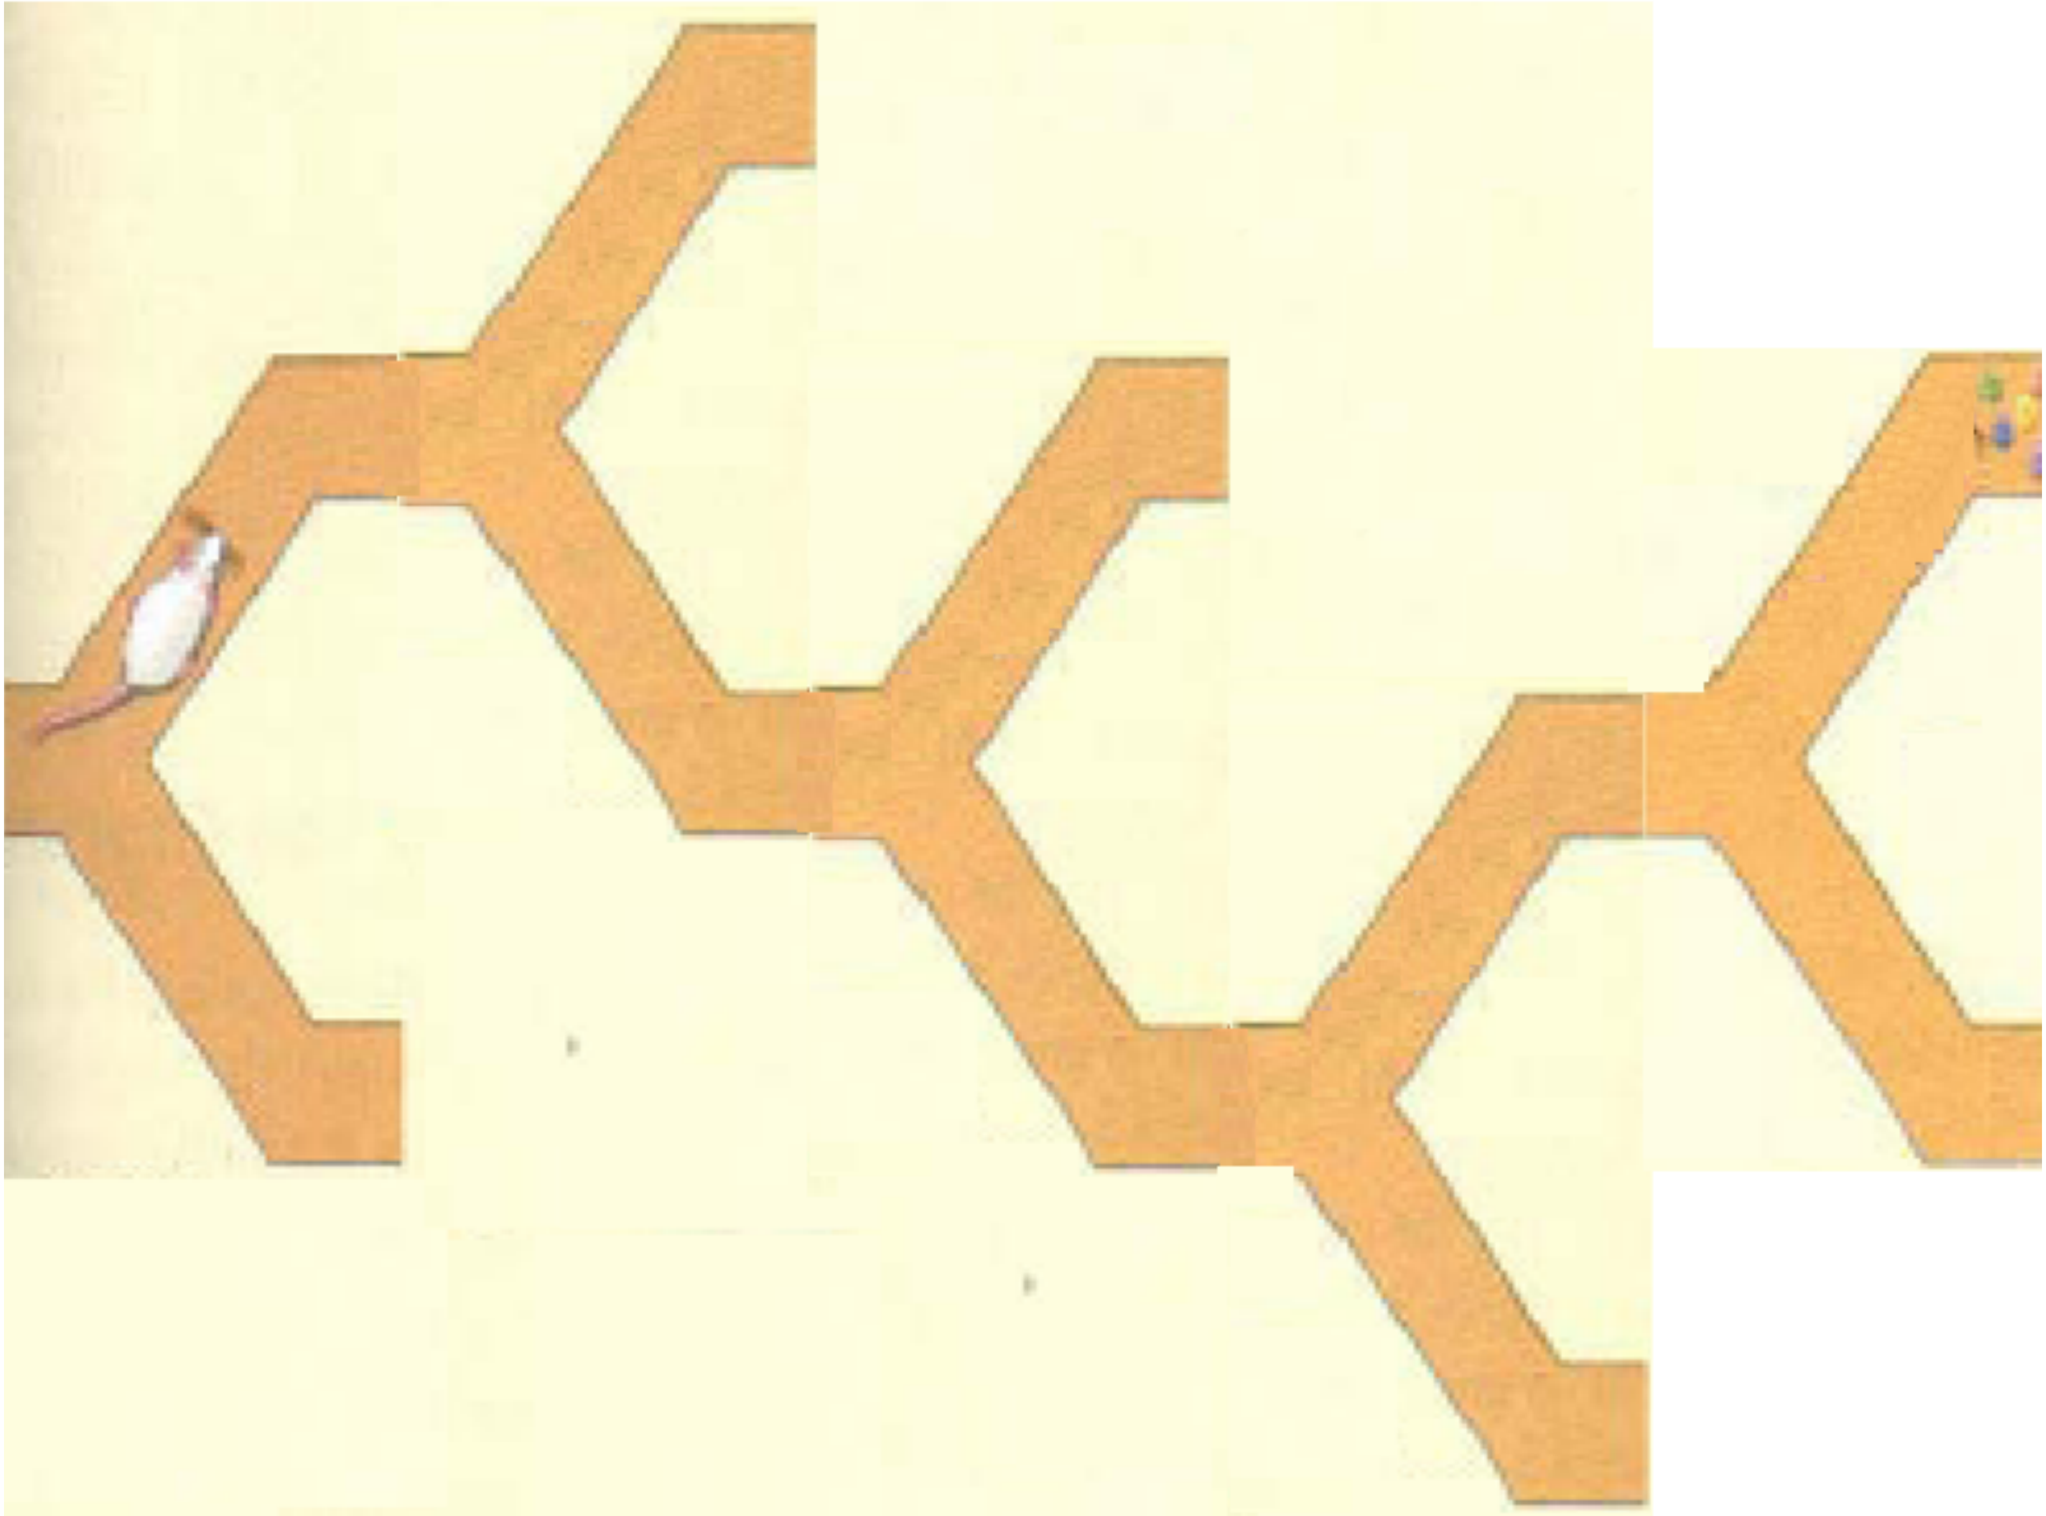

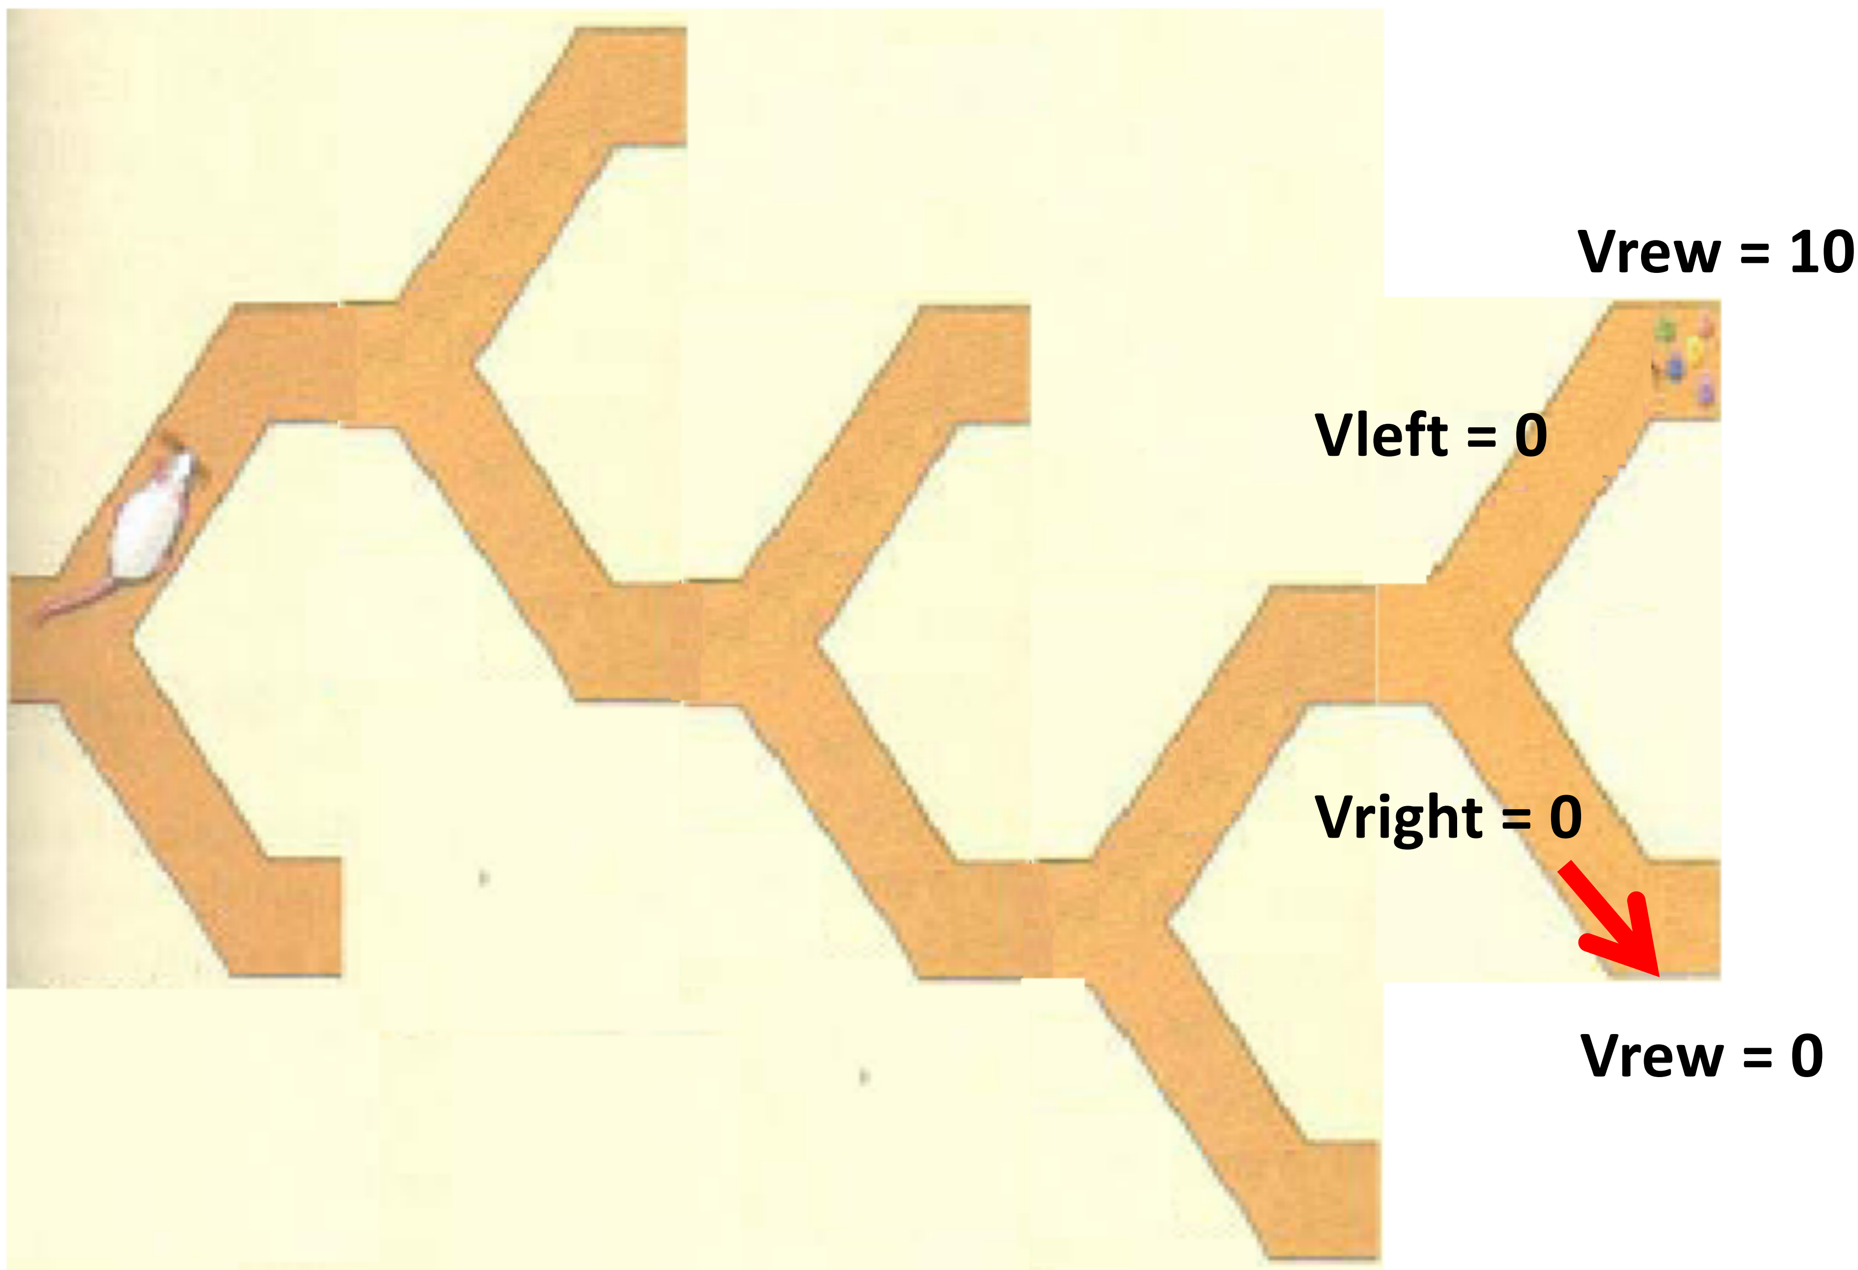
Vrew = 10
Vleft = 0
Vright = 0
Vrew = 0

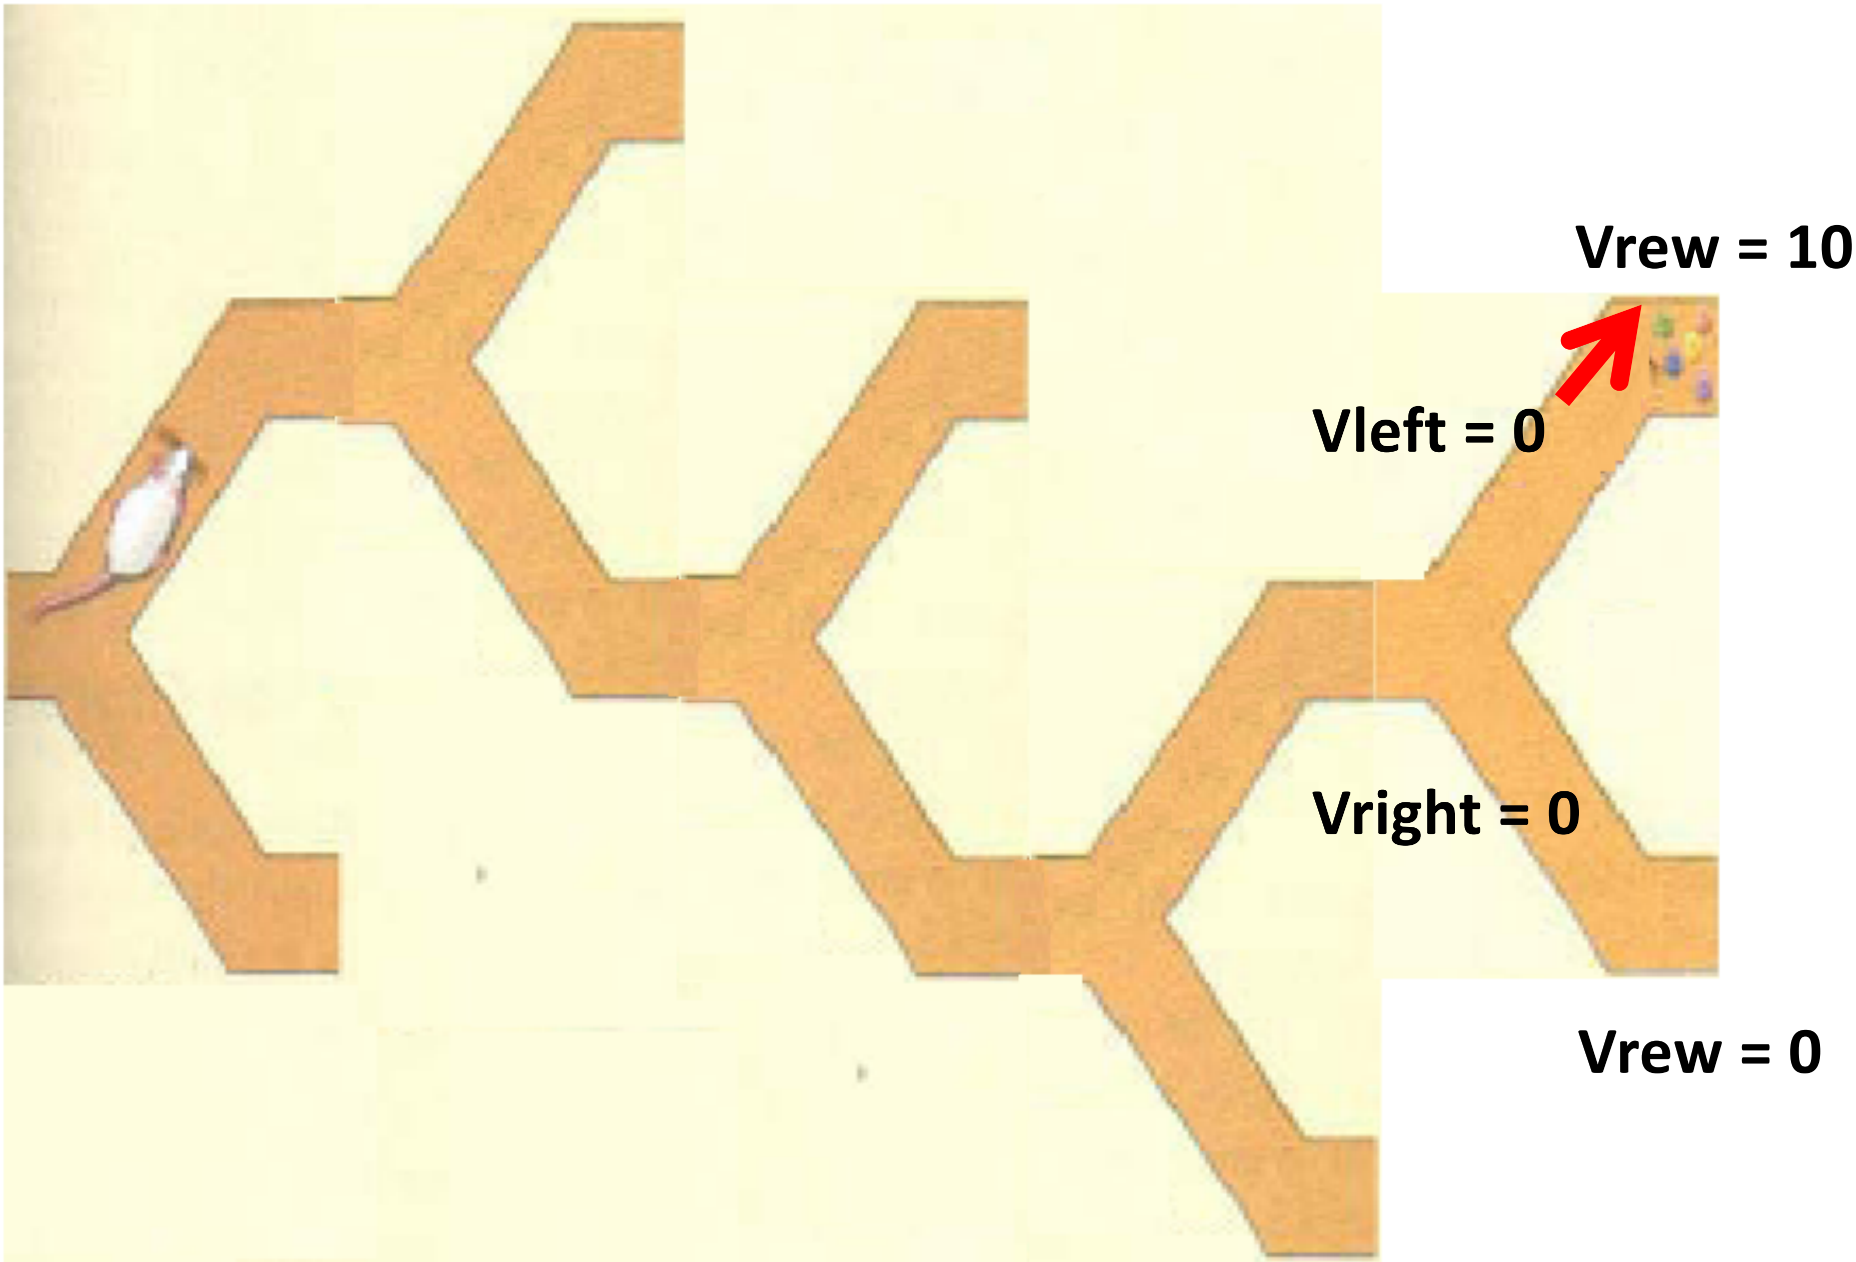

Vrew = 10
Vleft = 0
Vright = 0
Vrew = 0

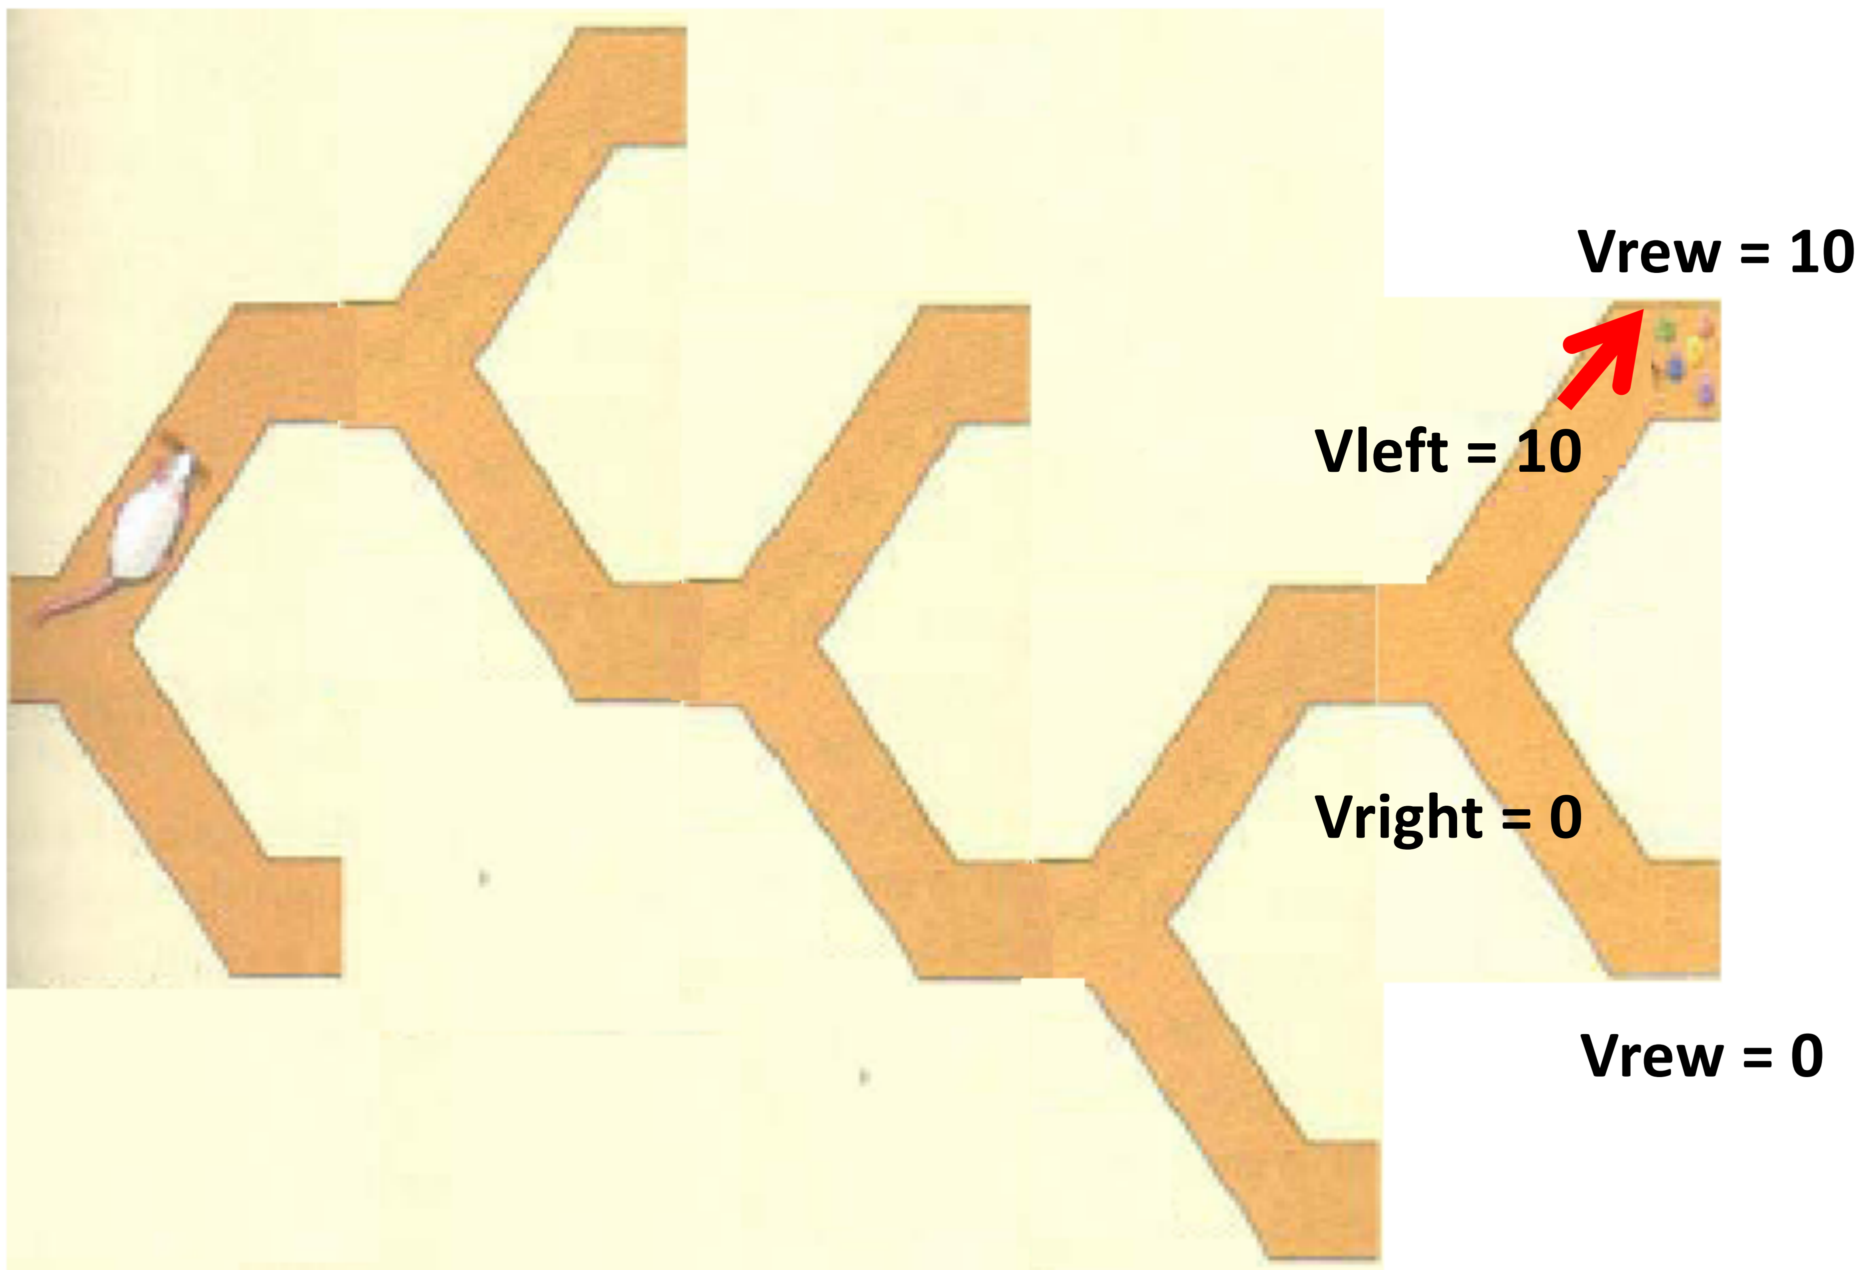
Vrew = 10
Vleft = 10
Vright = 0
Vrew = 0

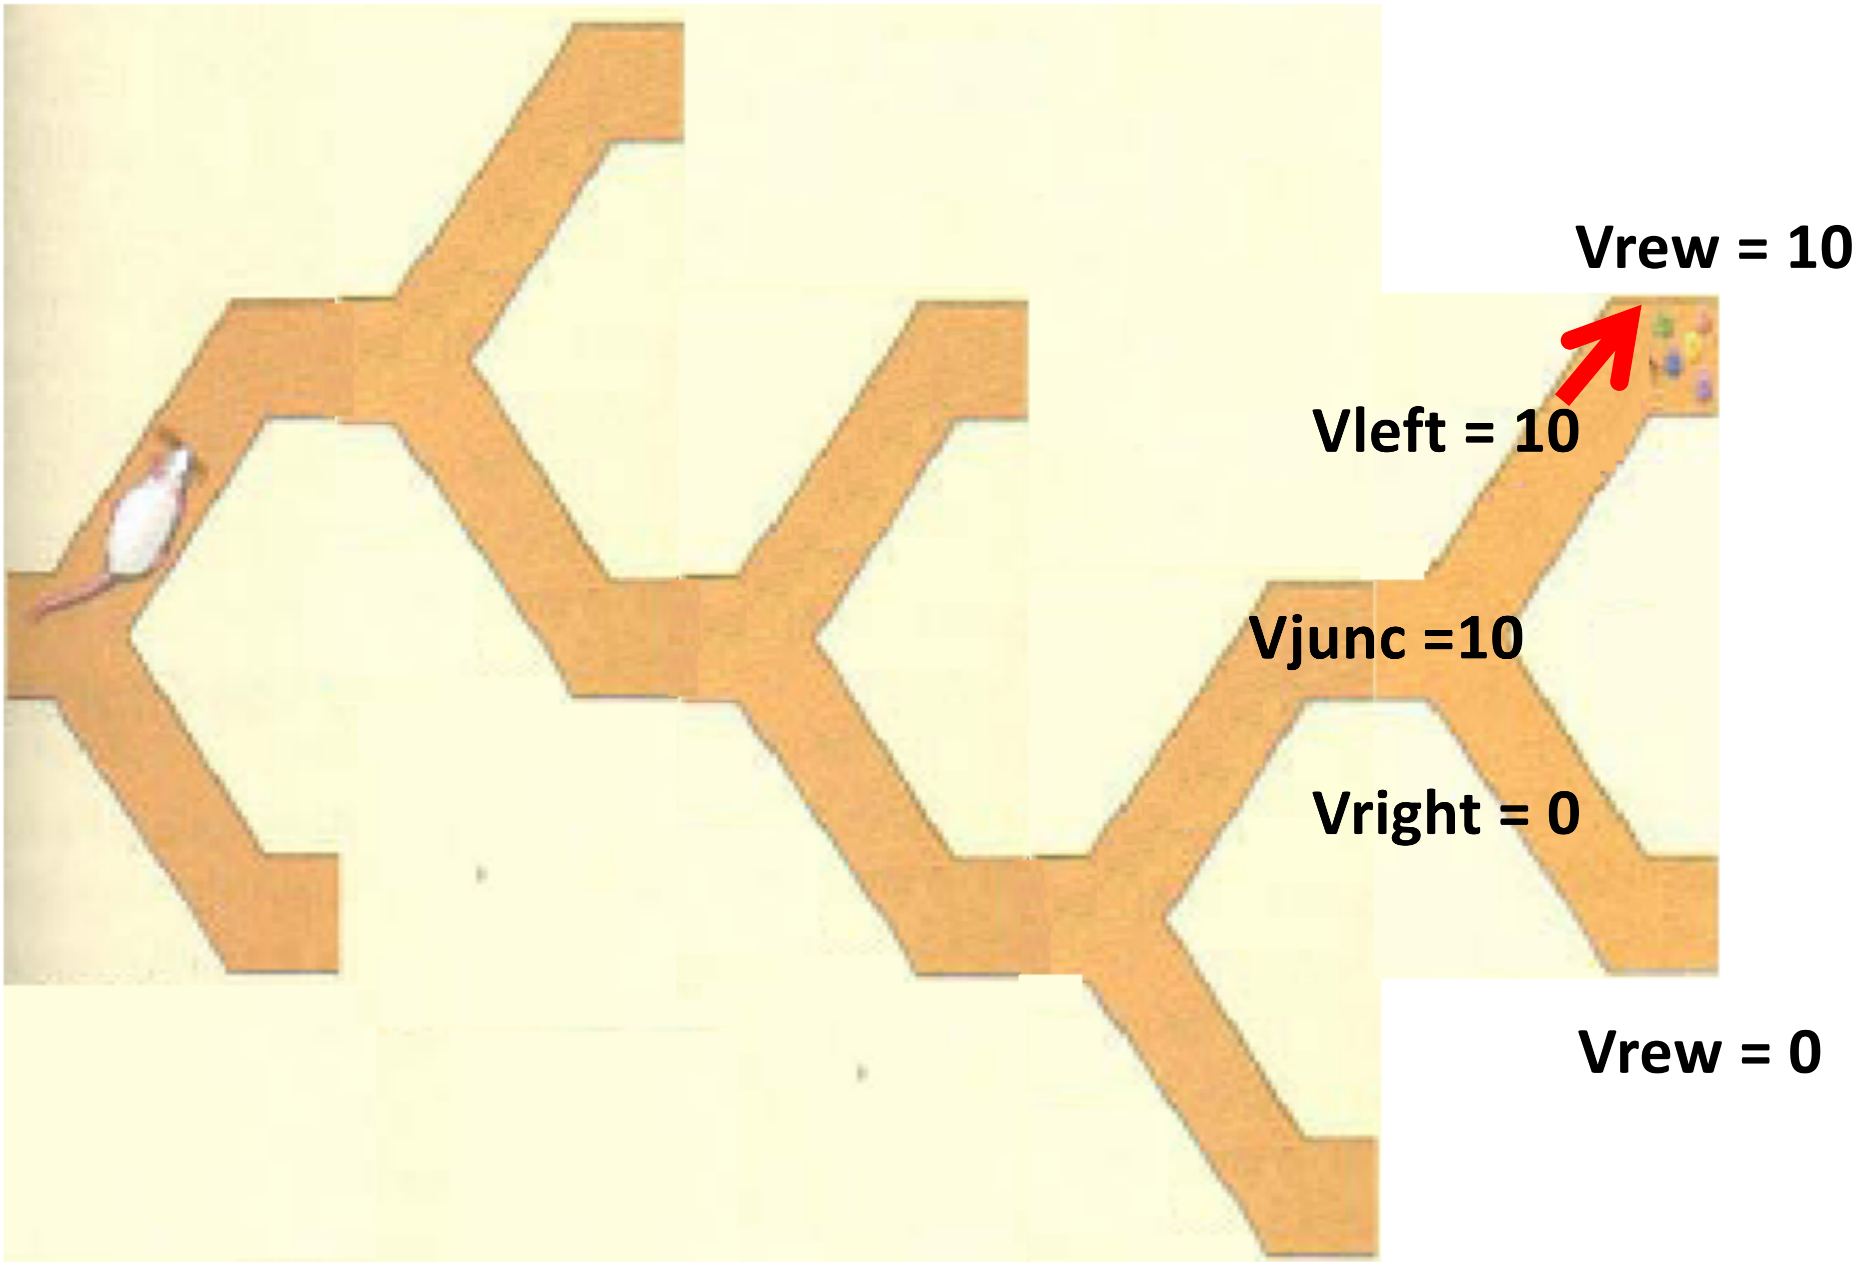
Vrew = 10
Vleft = 10
Vjunc =10
Vright = 0
Vrew = 0

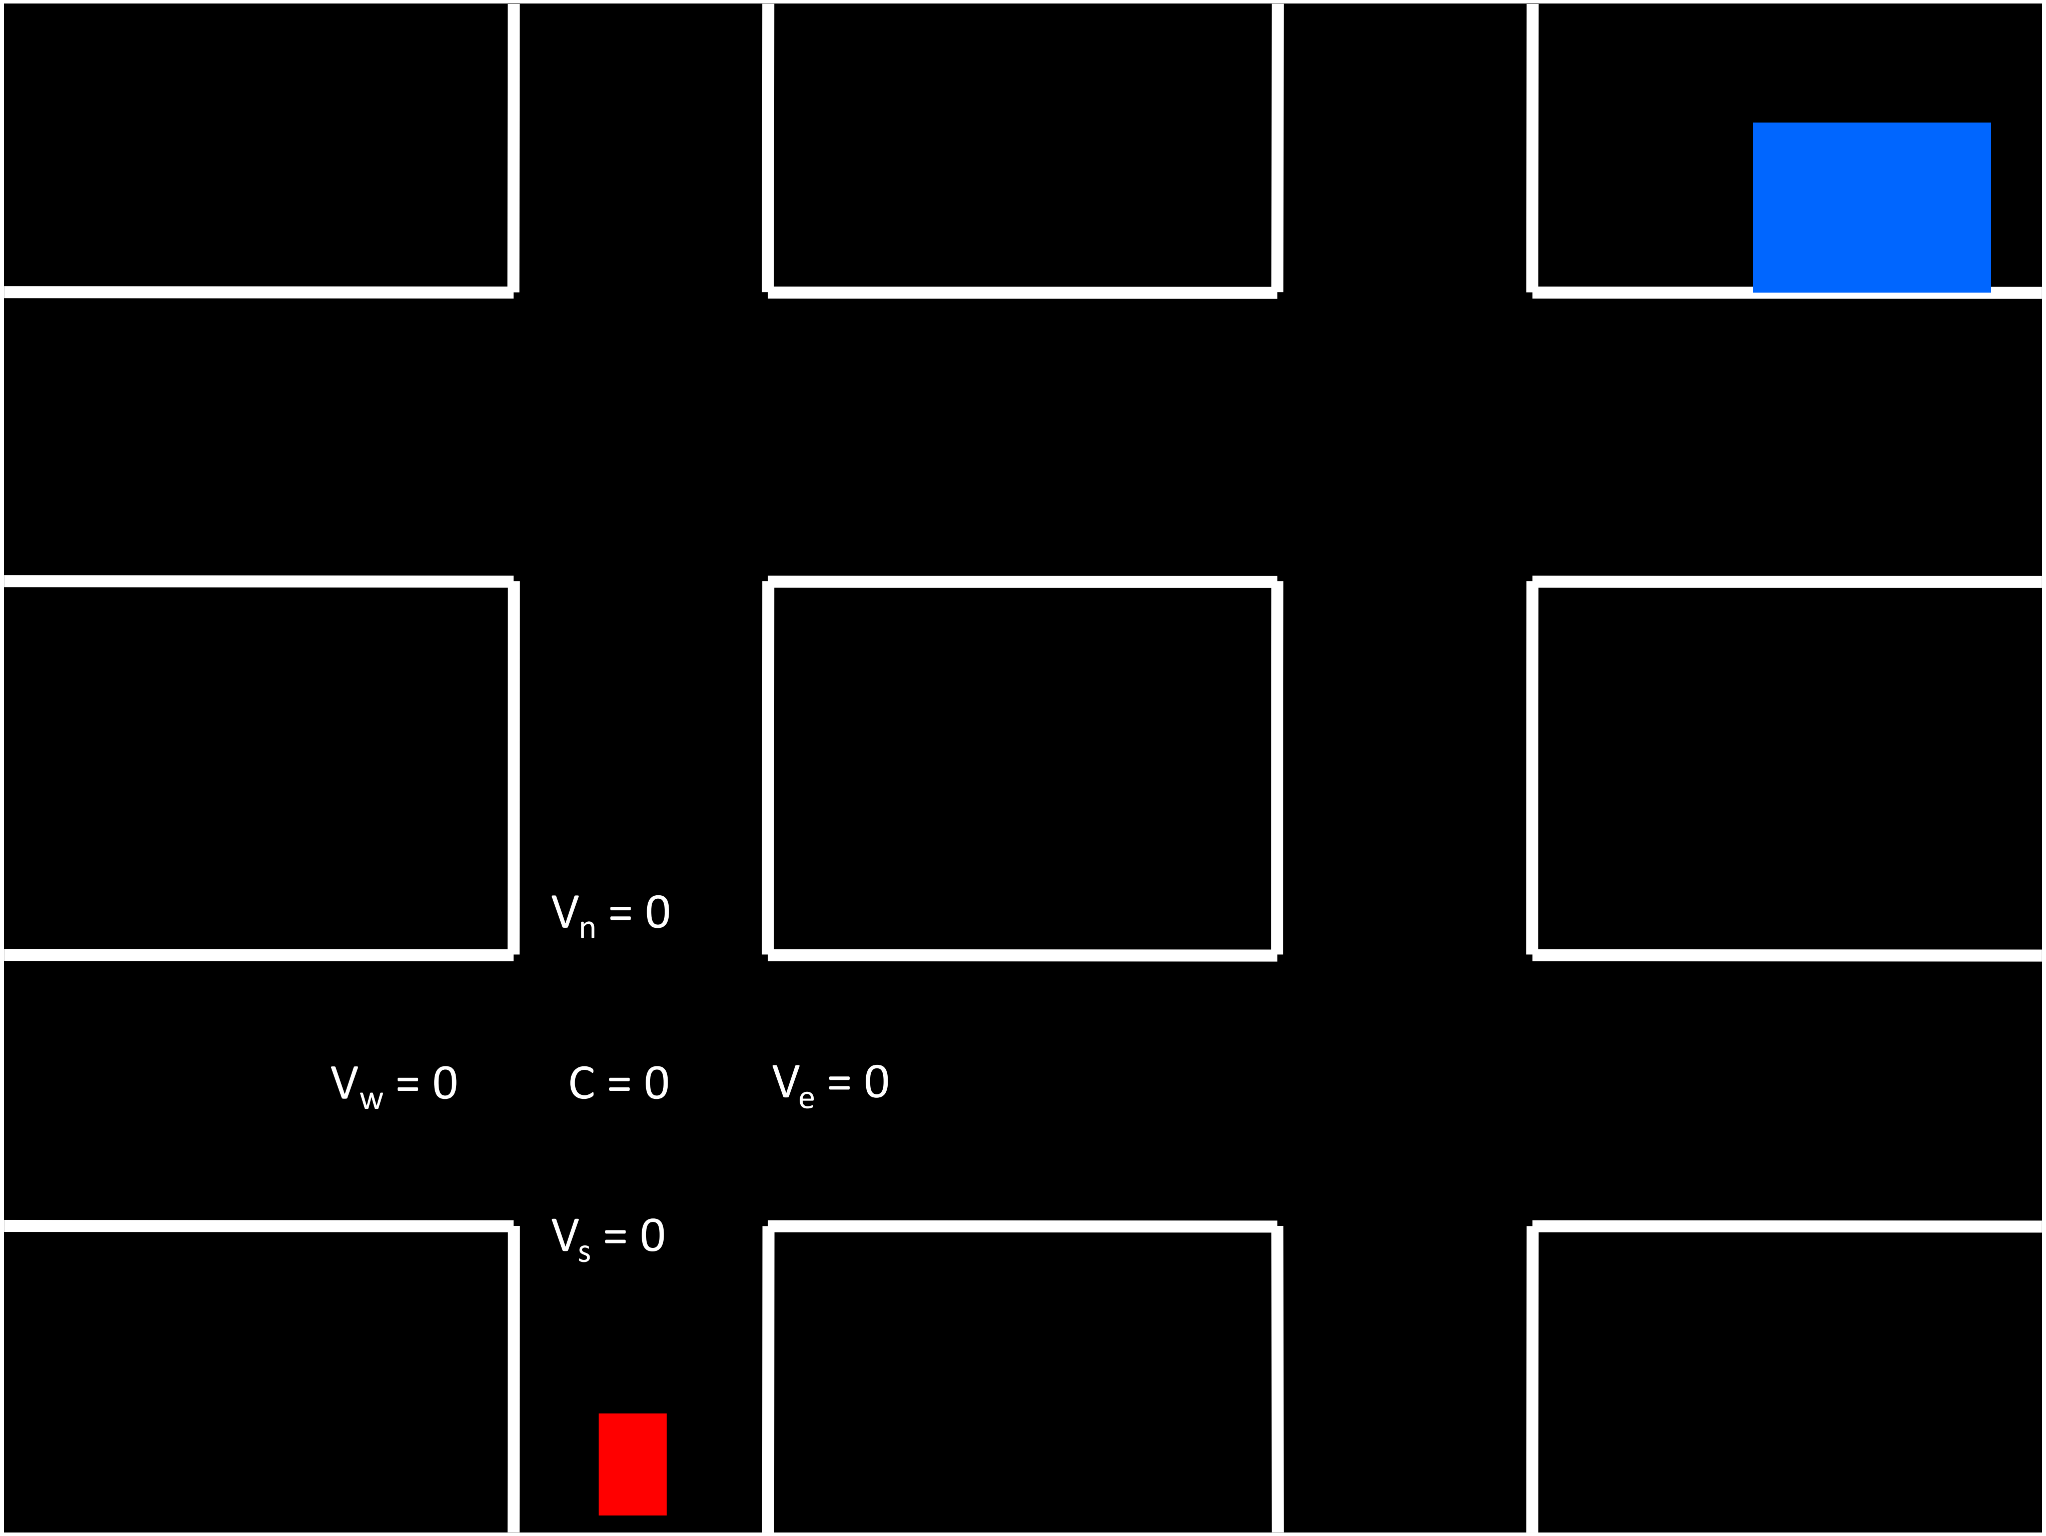
$V_n = 0$
$V_w = 0$
$C = 0$
$V_e = 0$
$V_s = 0$

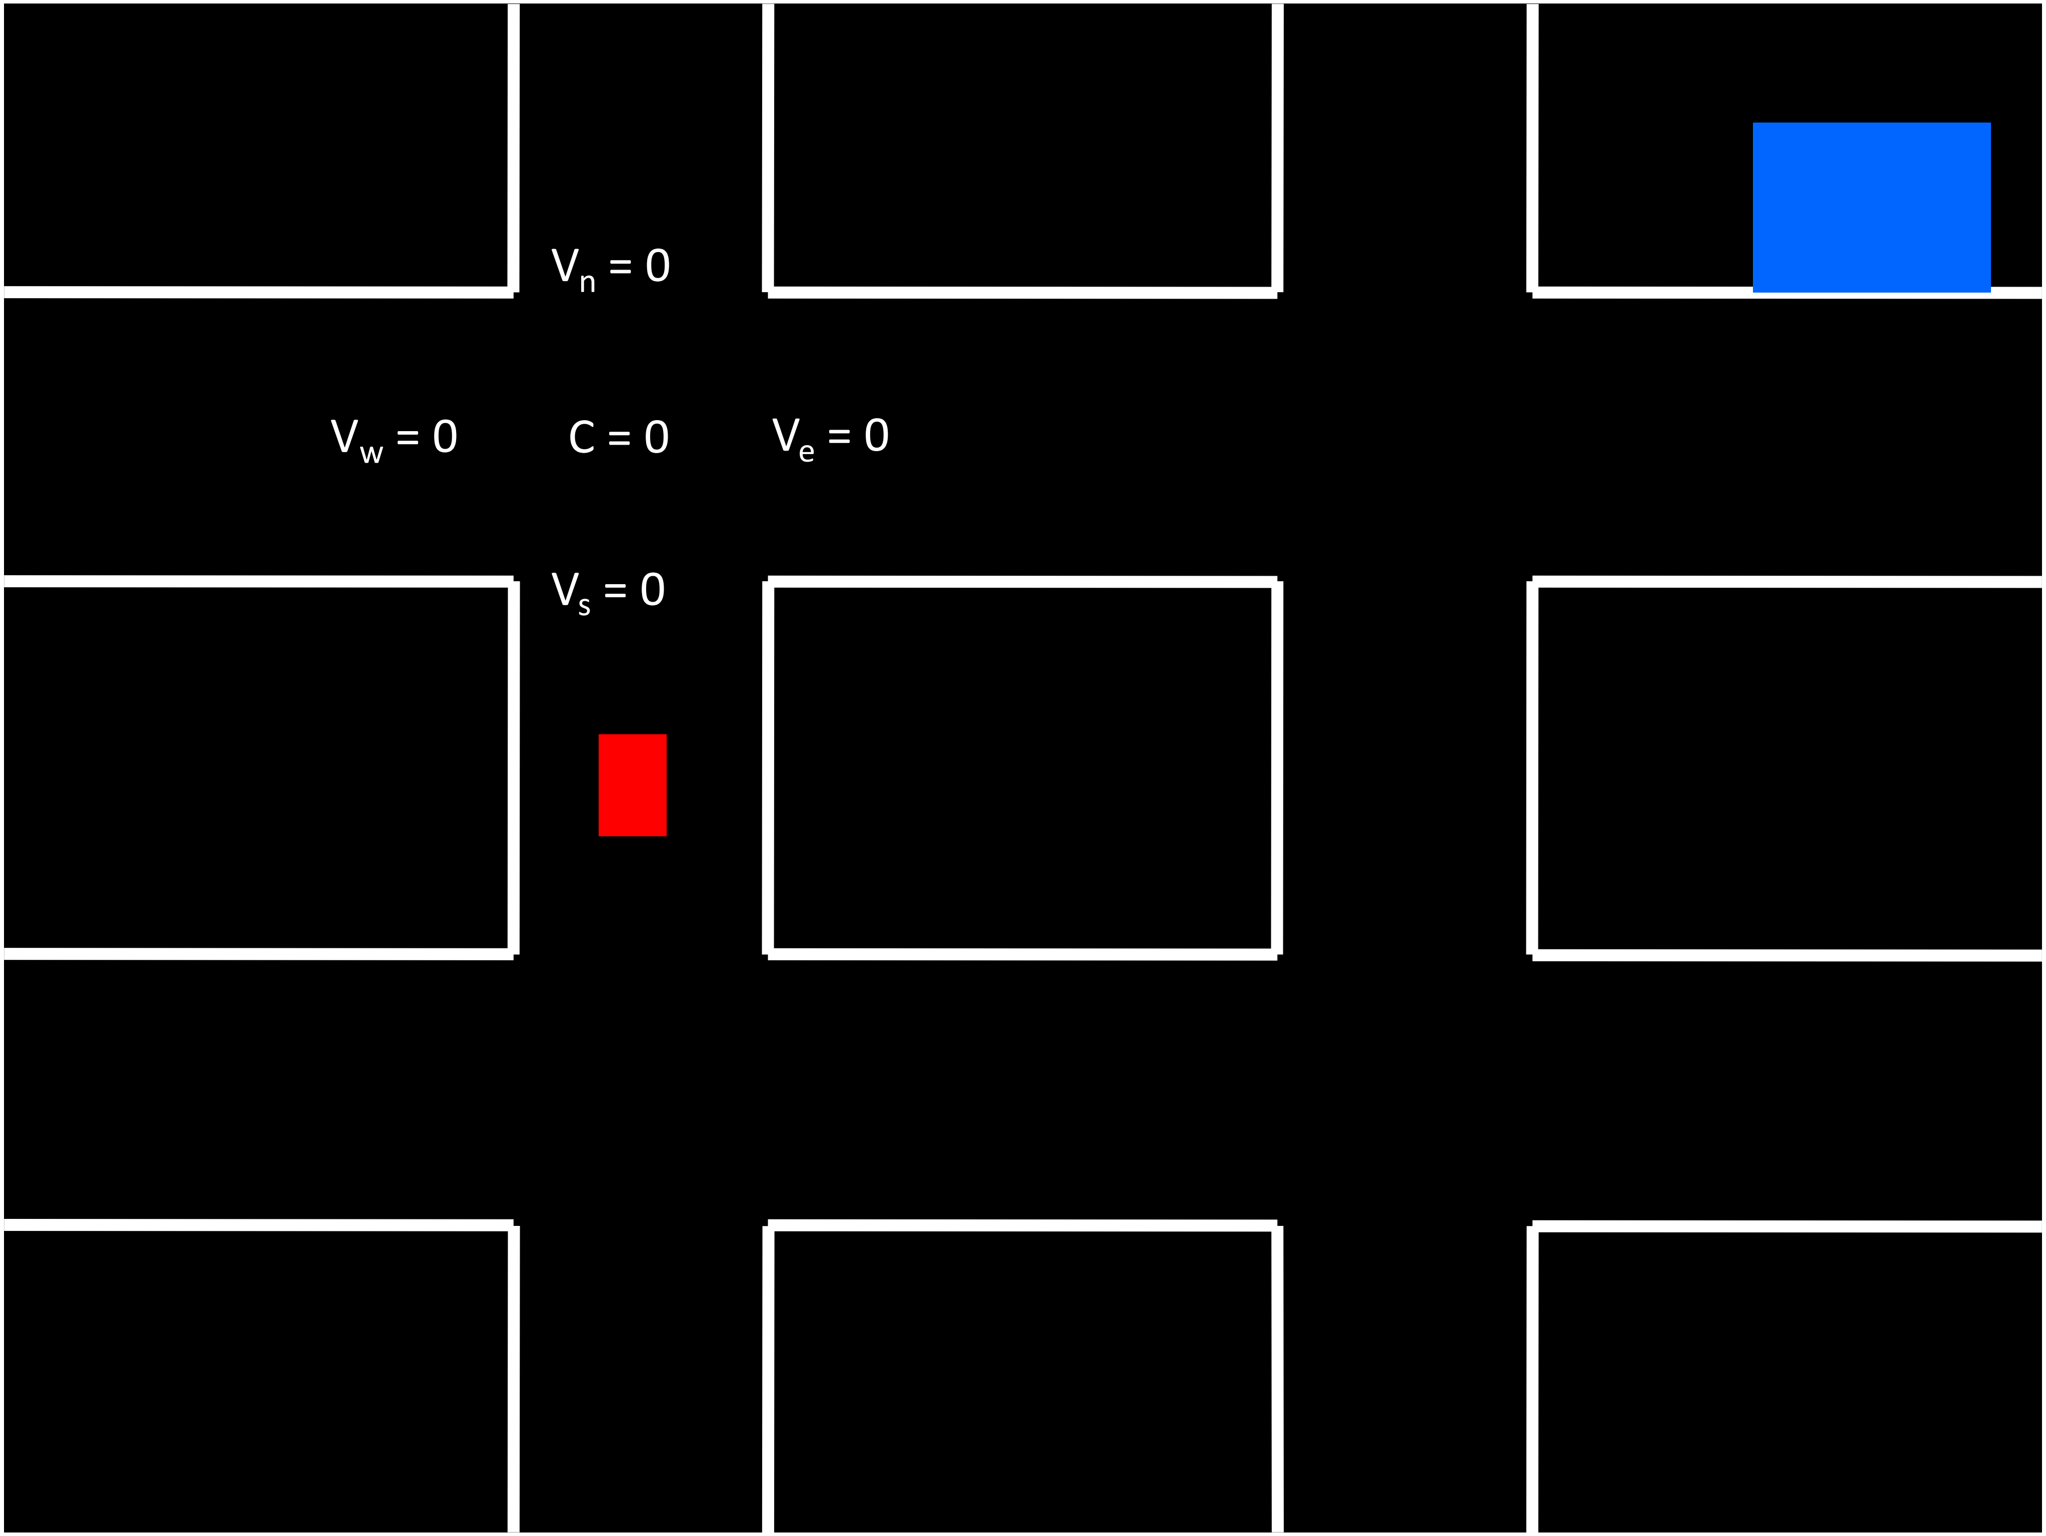
$V_n = 0$
$V_w = 0$
$C = 0$
$V_e = 0$
$V_s = 0$

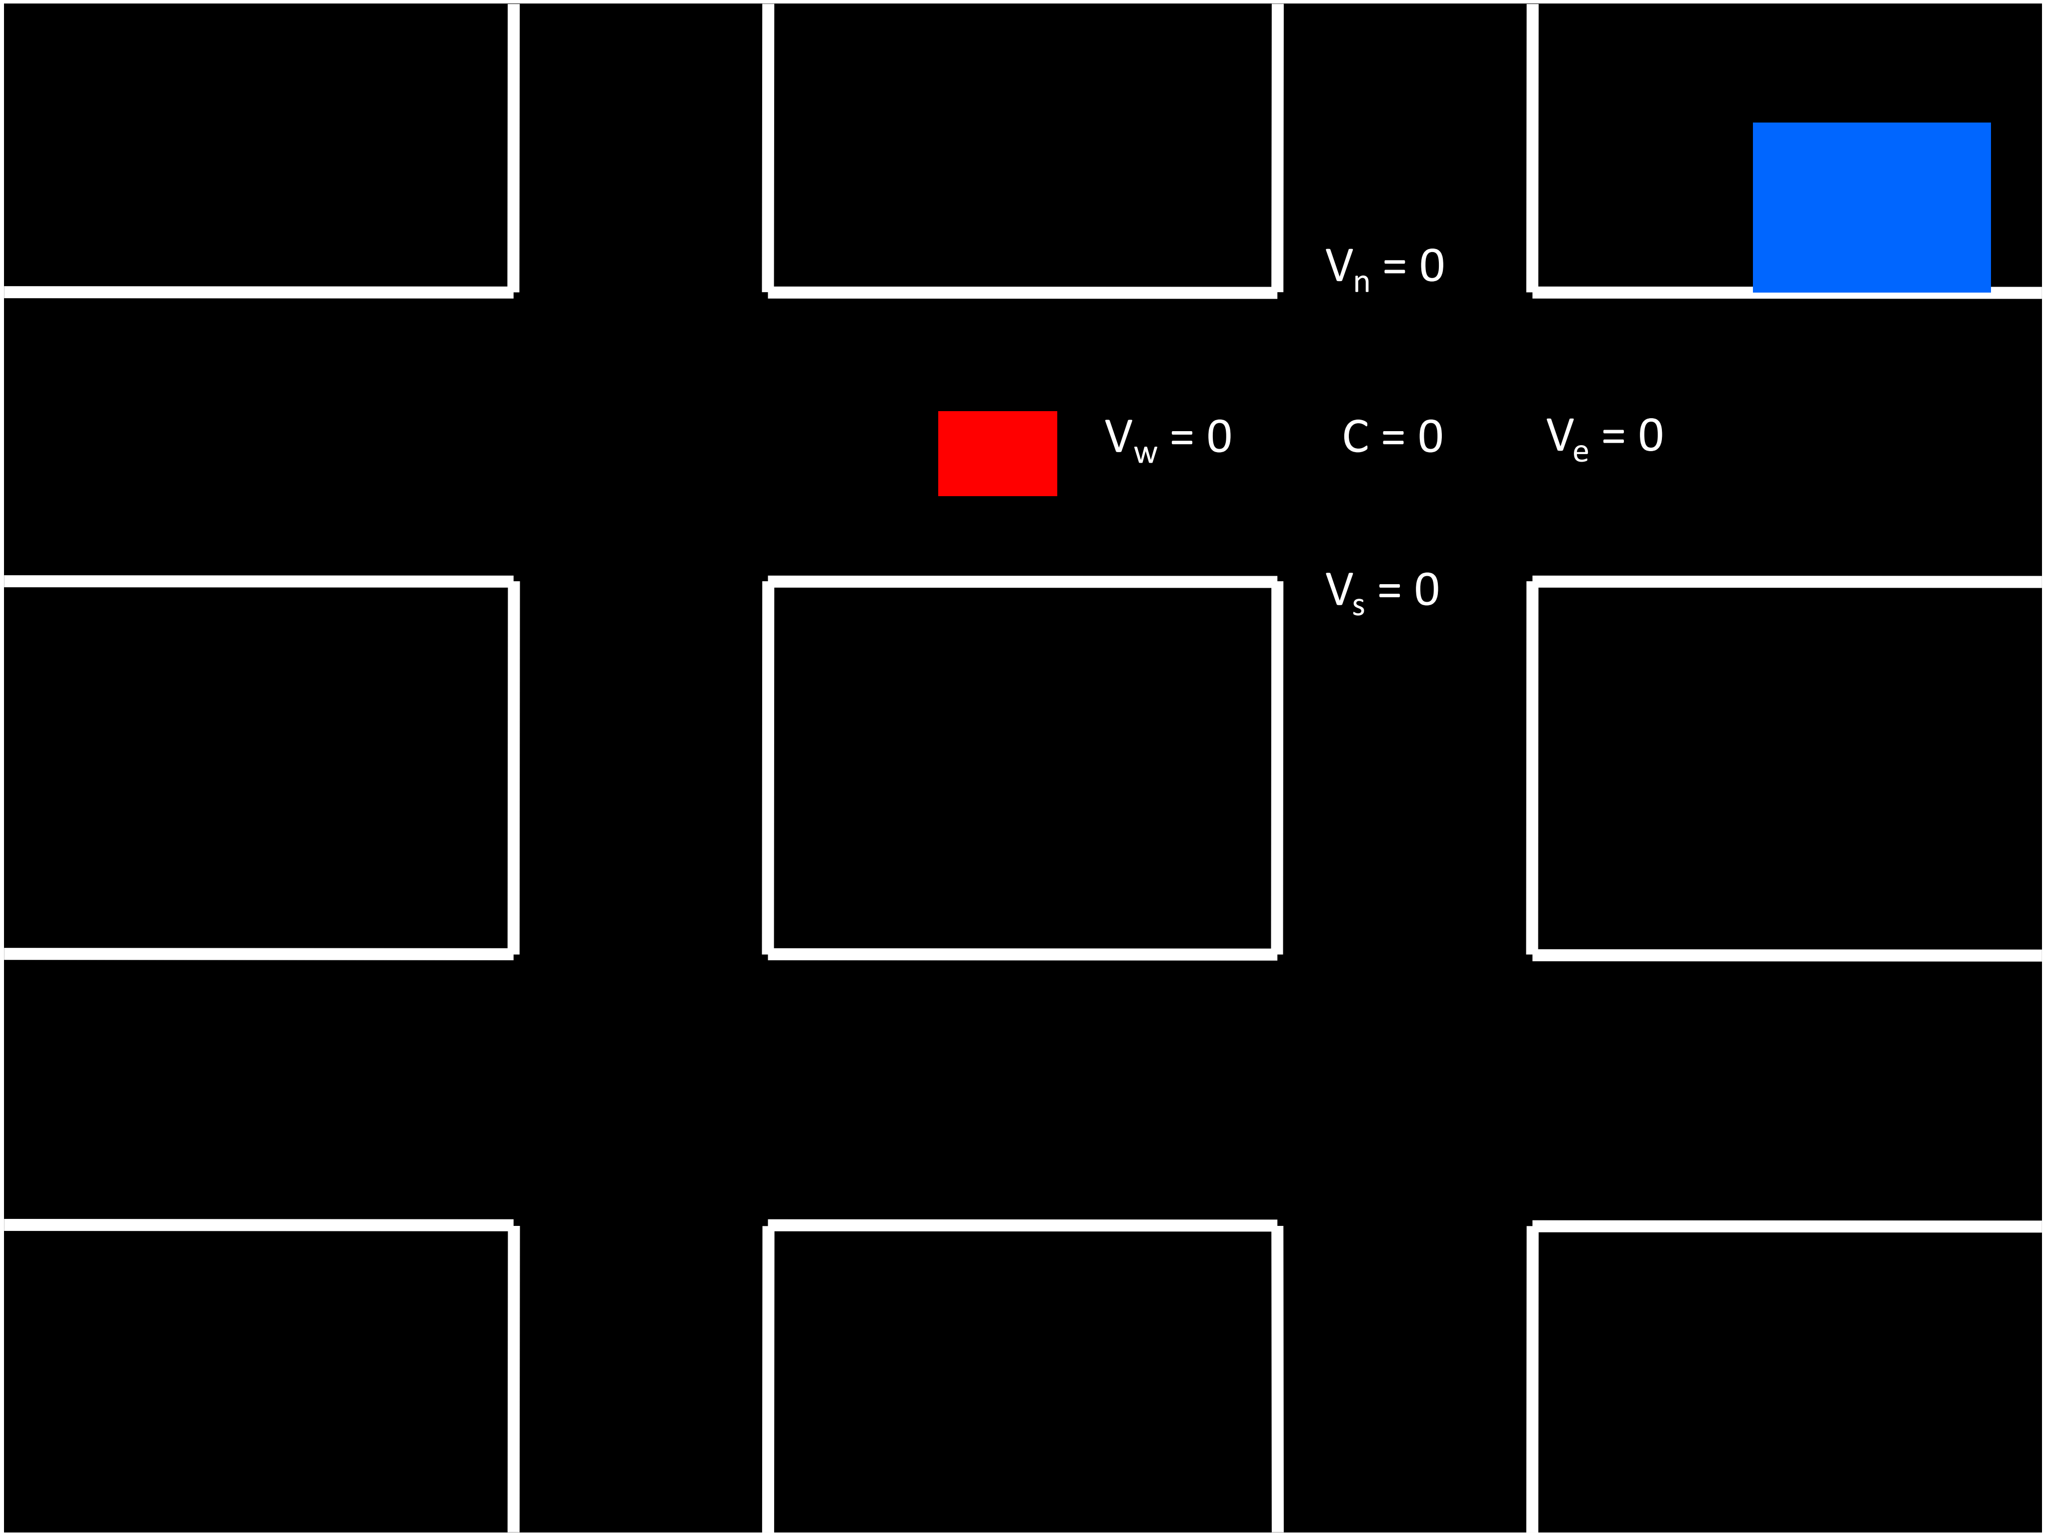

$V_n = 0$
$V_w = 0$
$C = 0$
$V_e = 0$
$V_s = 0$

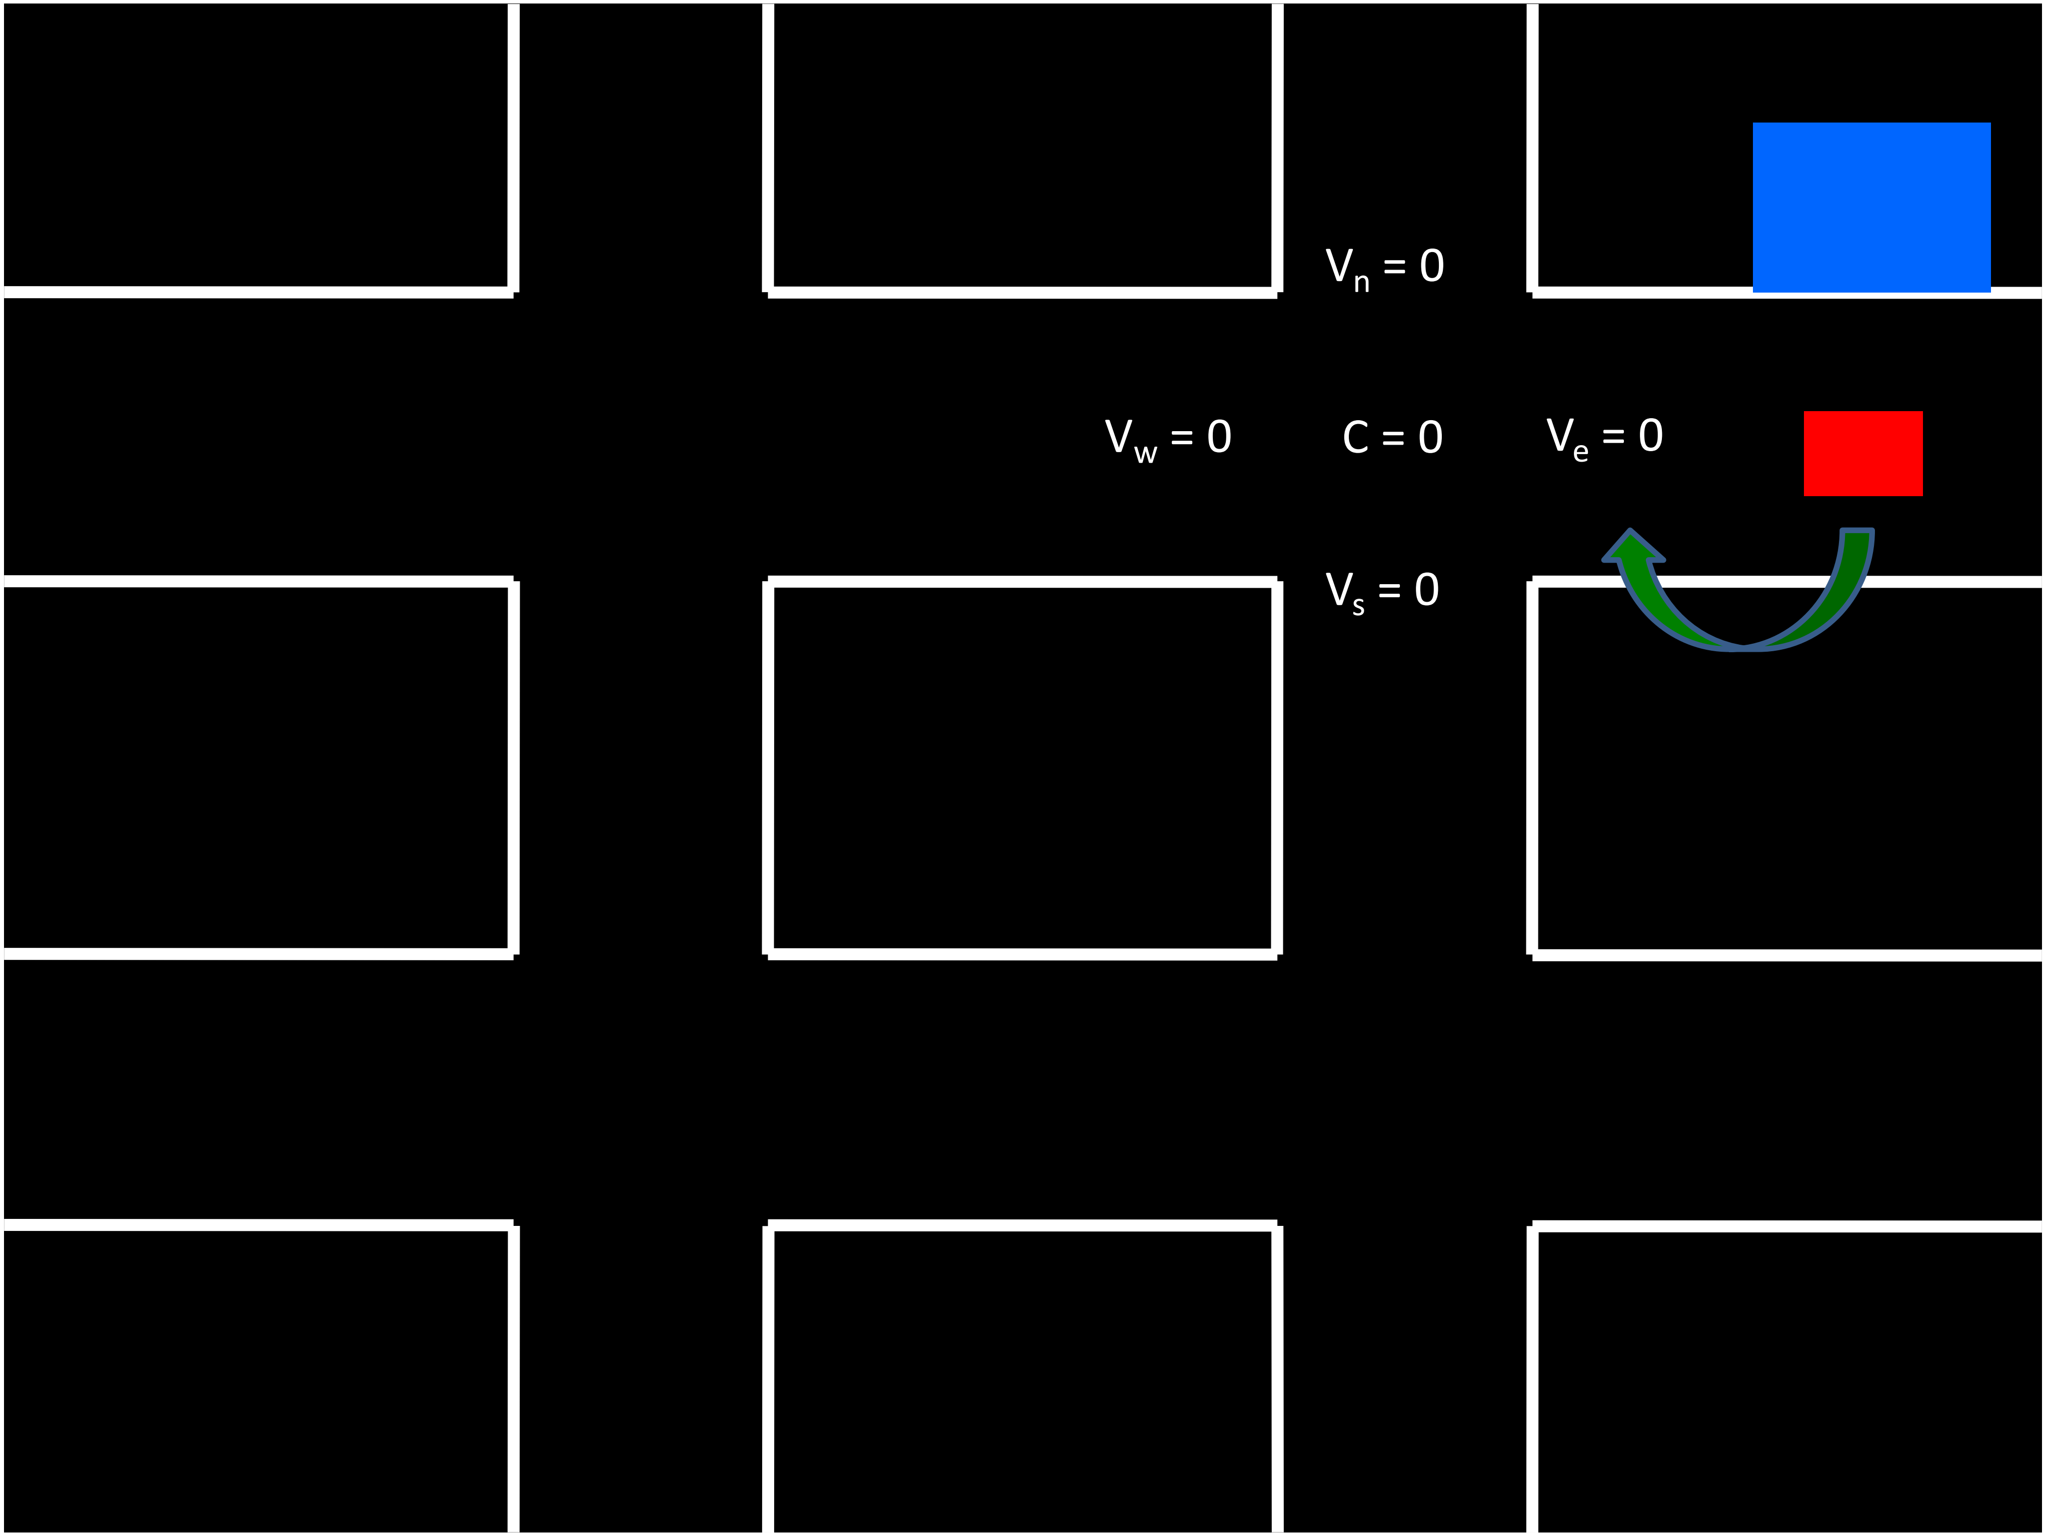

$V_n = 0$
$V_w = 0$
$C = 0$
$V_e = 0$
$V_s = 0$

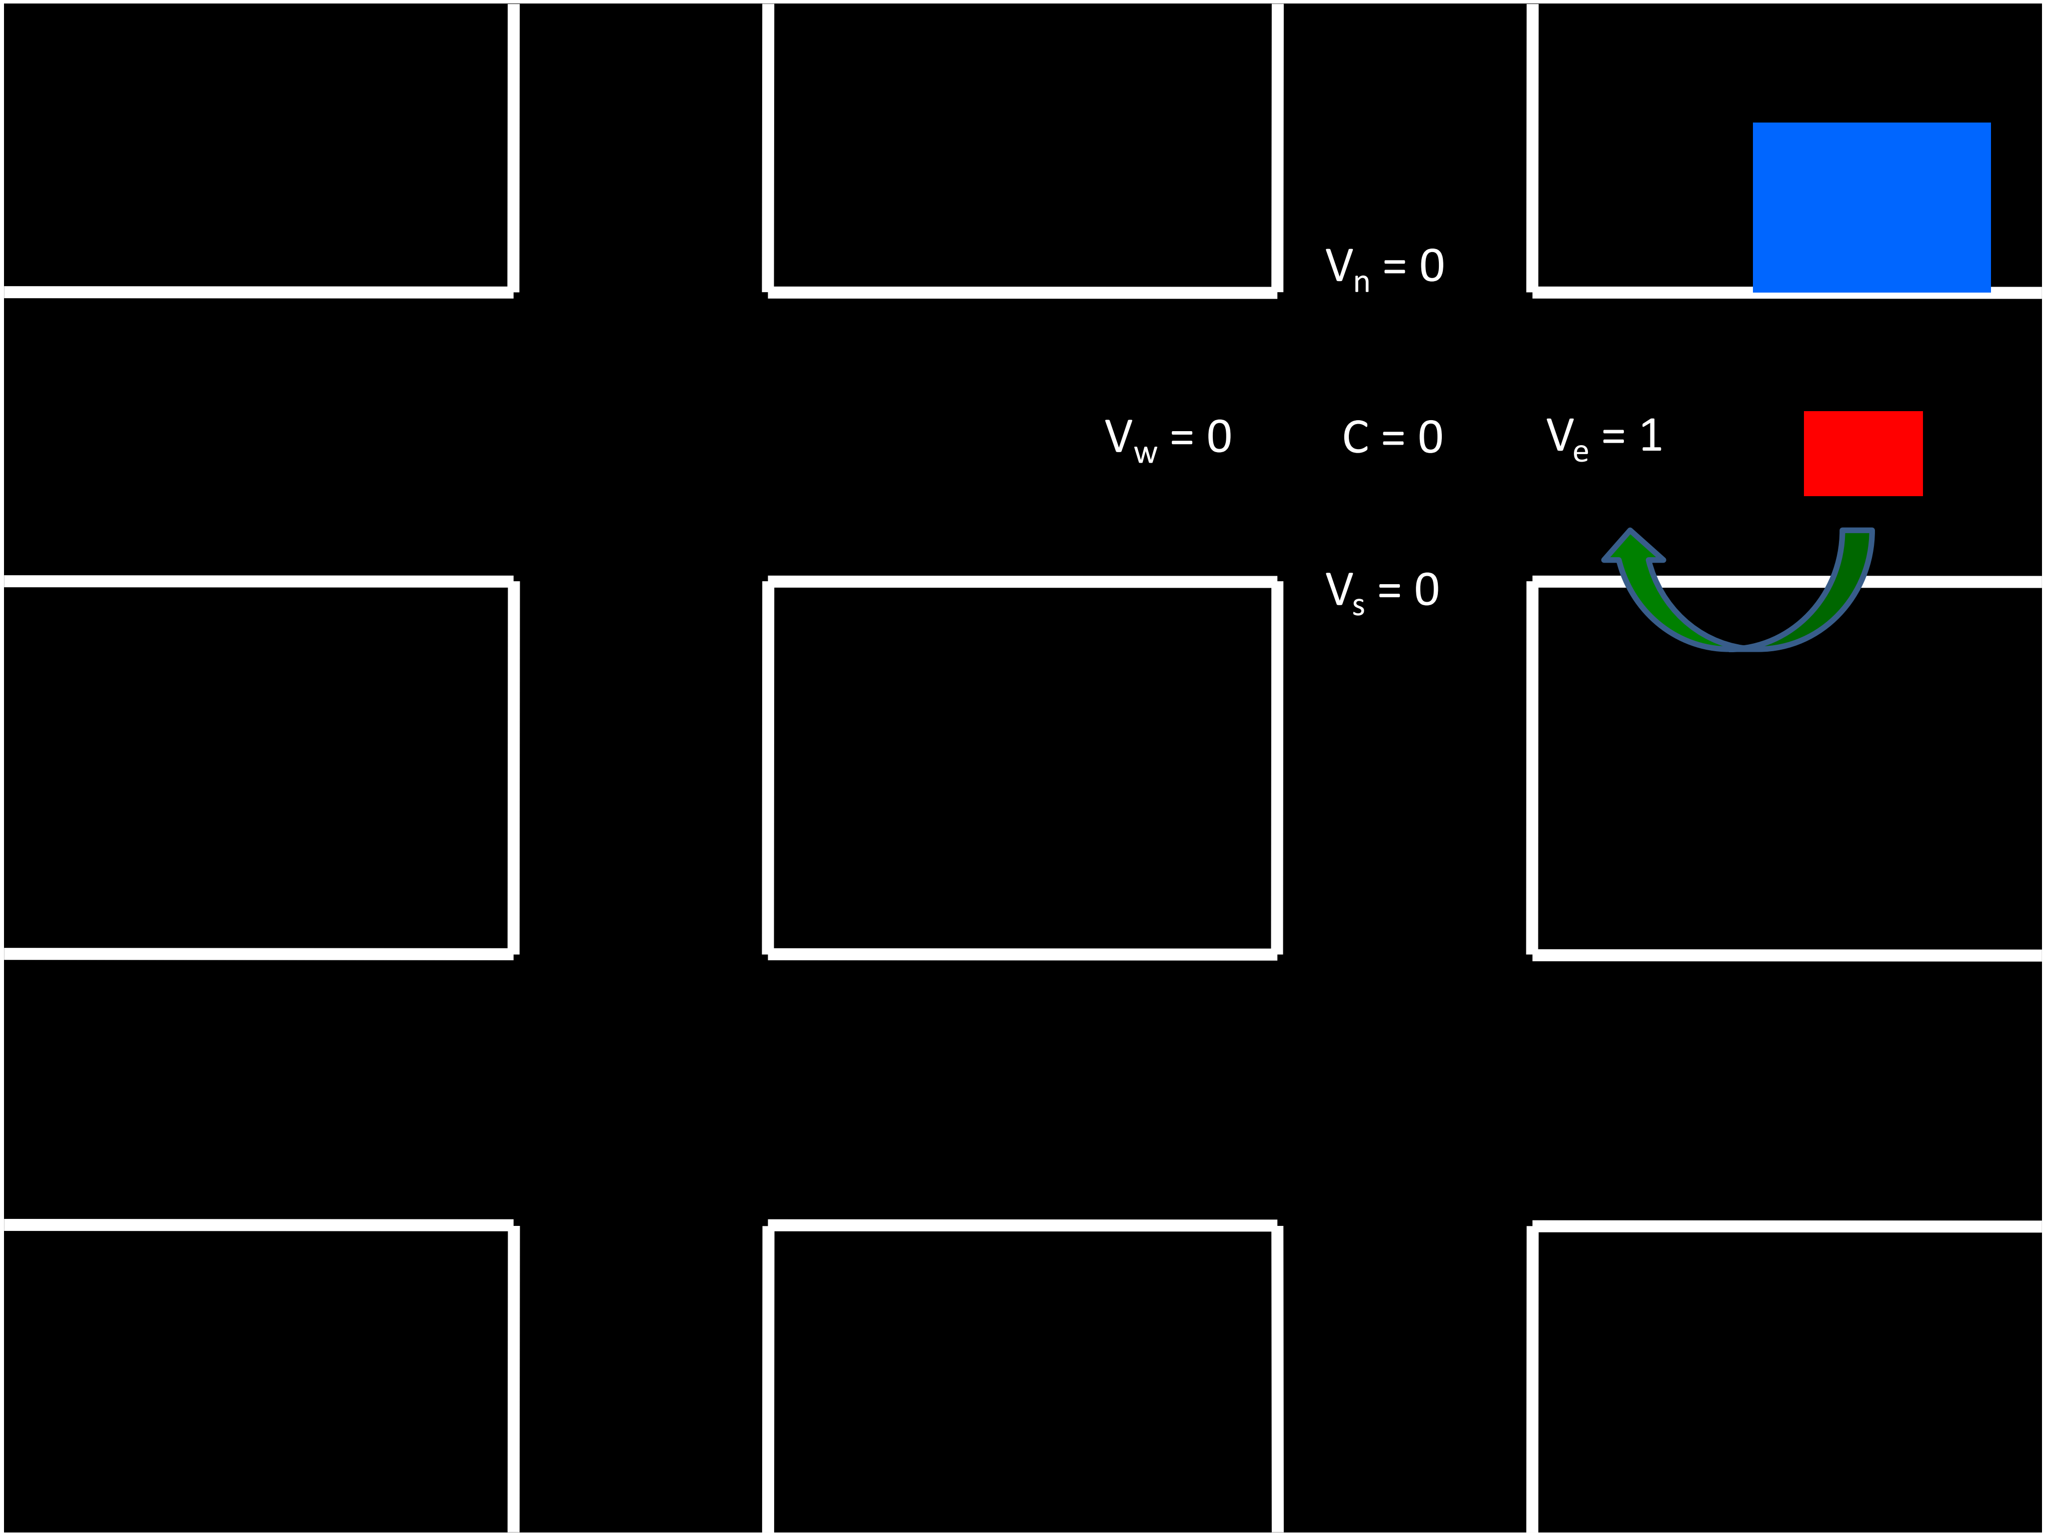

$V_n = 0$
$V_w = 0$
$C = 0$
$V_e = 1$
$V_s = 0$

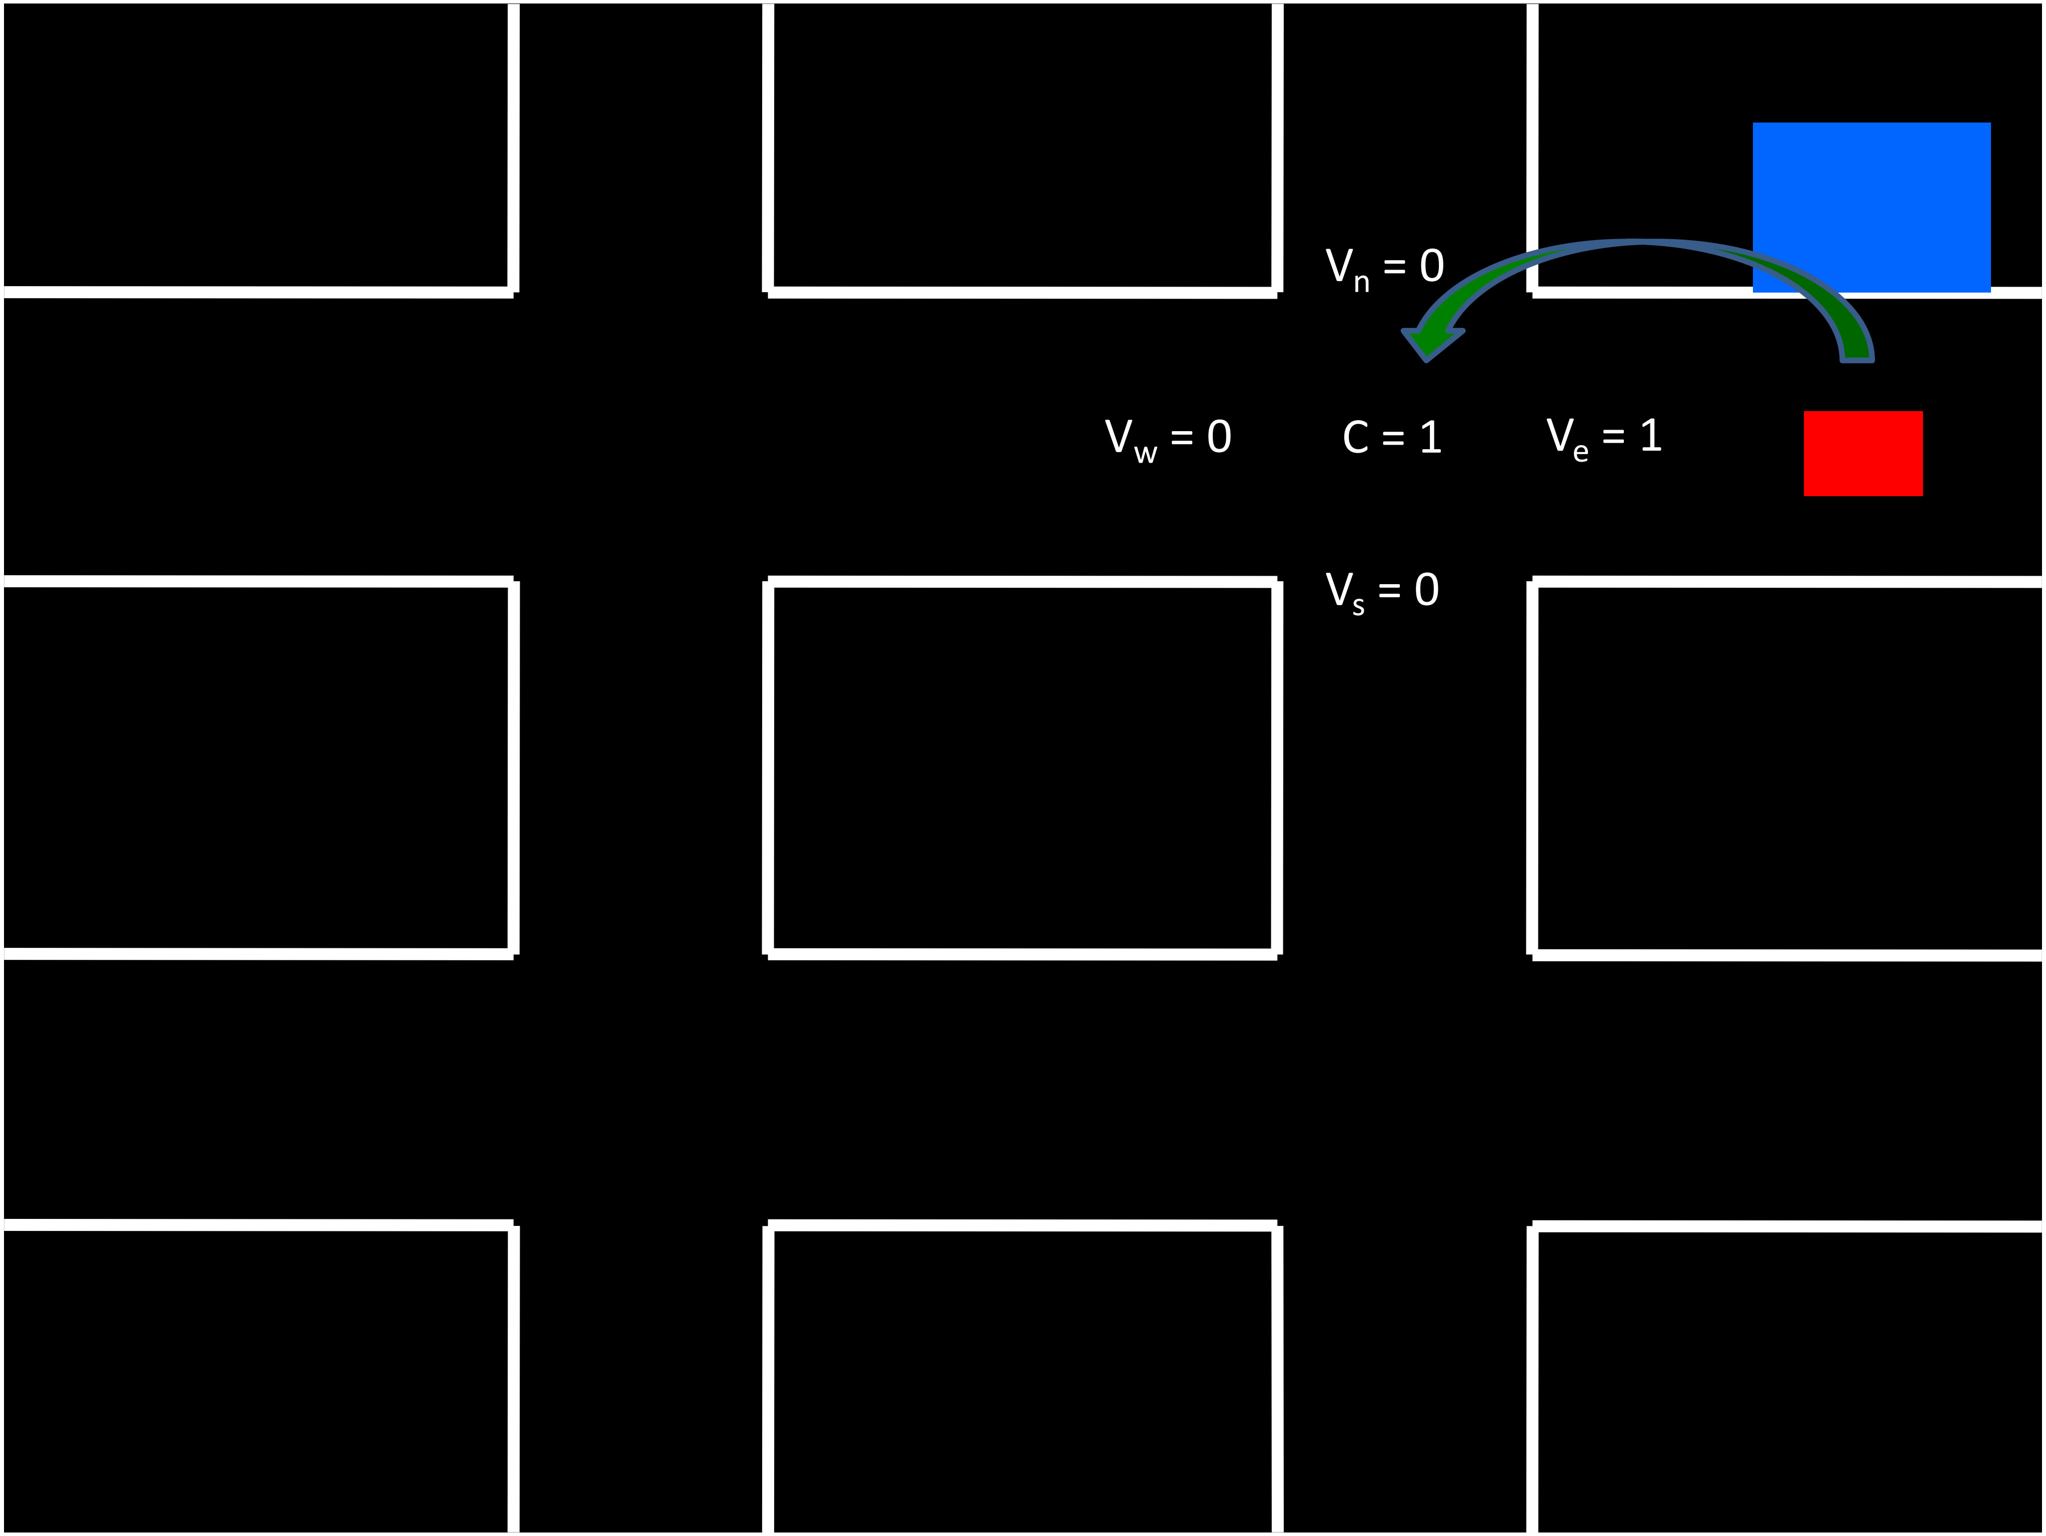

$V_n = 0$
$V_w = 0$
$C = 1$
$V_e = 1$
$V_s = 0$

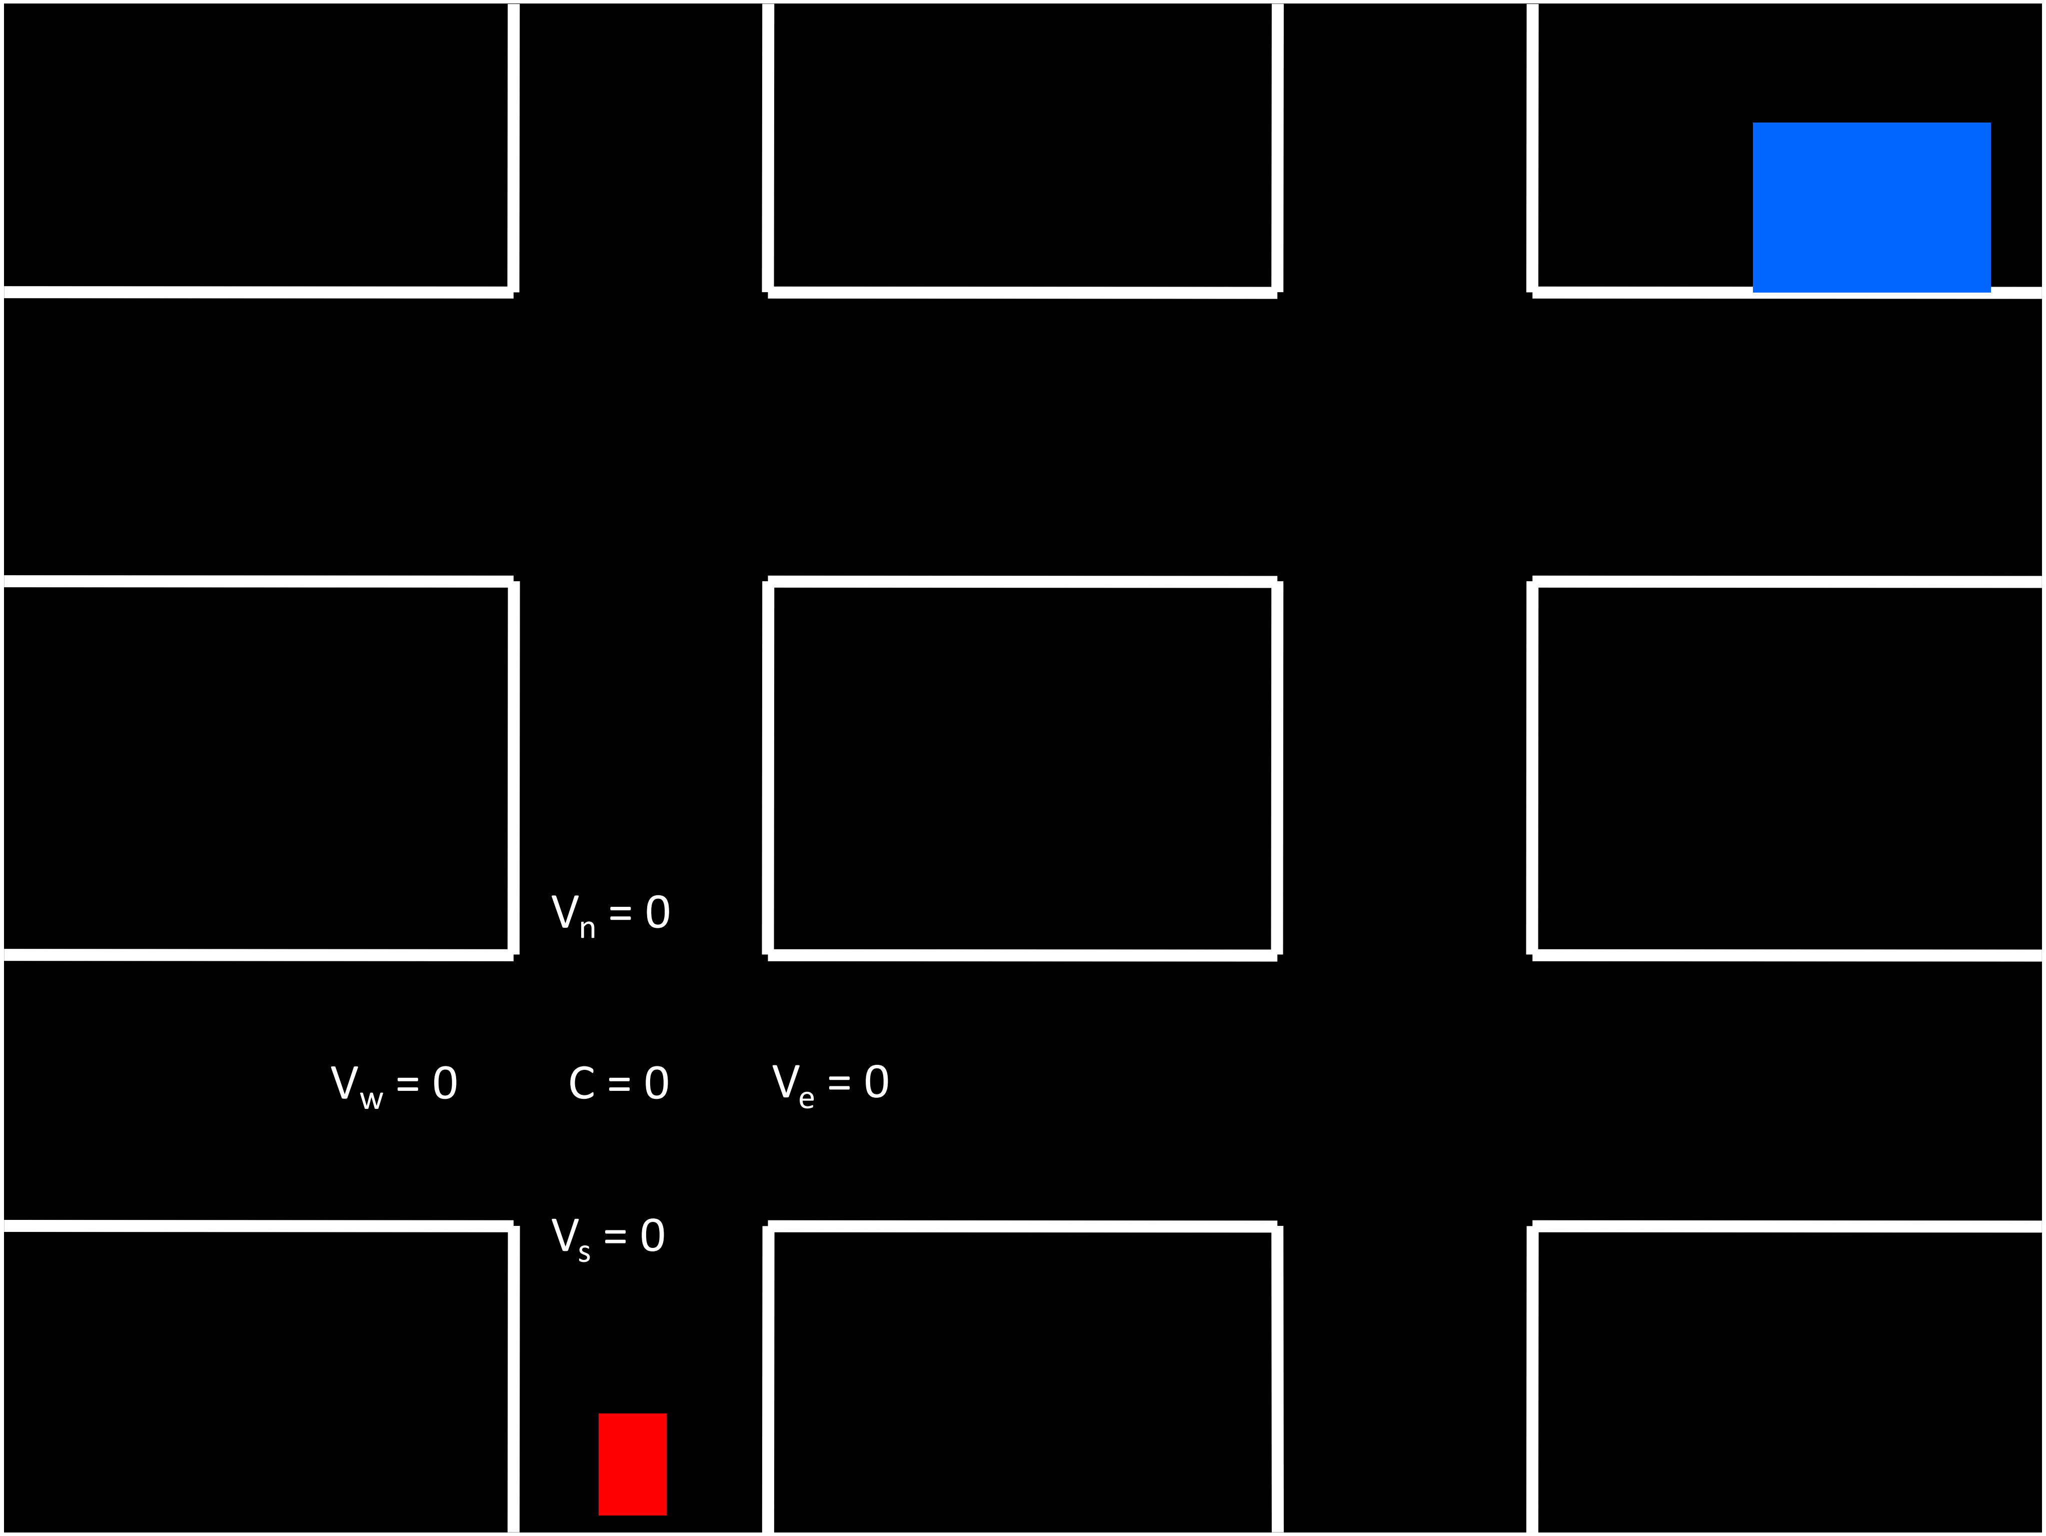
$V_n = 0$
$V_w = 0$
$C = 0$
$V_e = 0$
$V_s = 0$

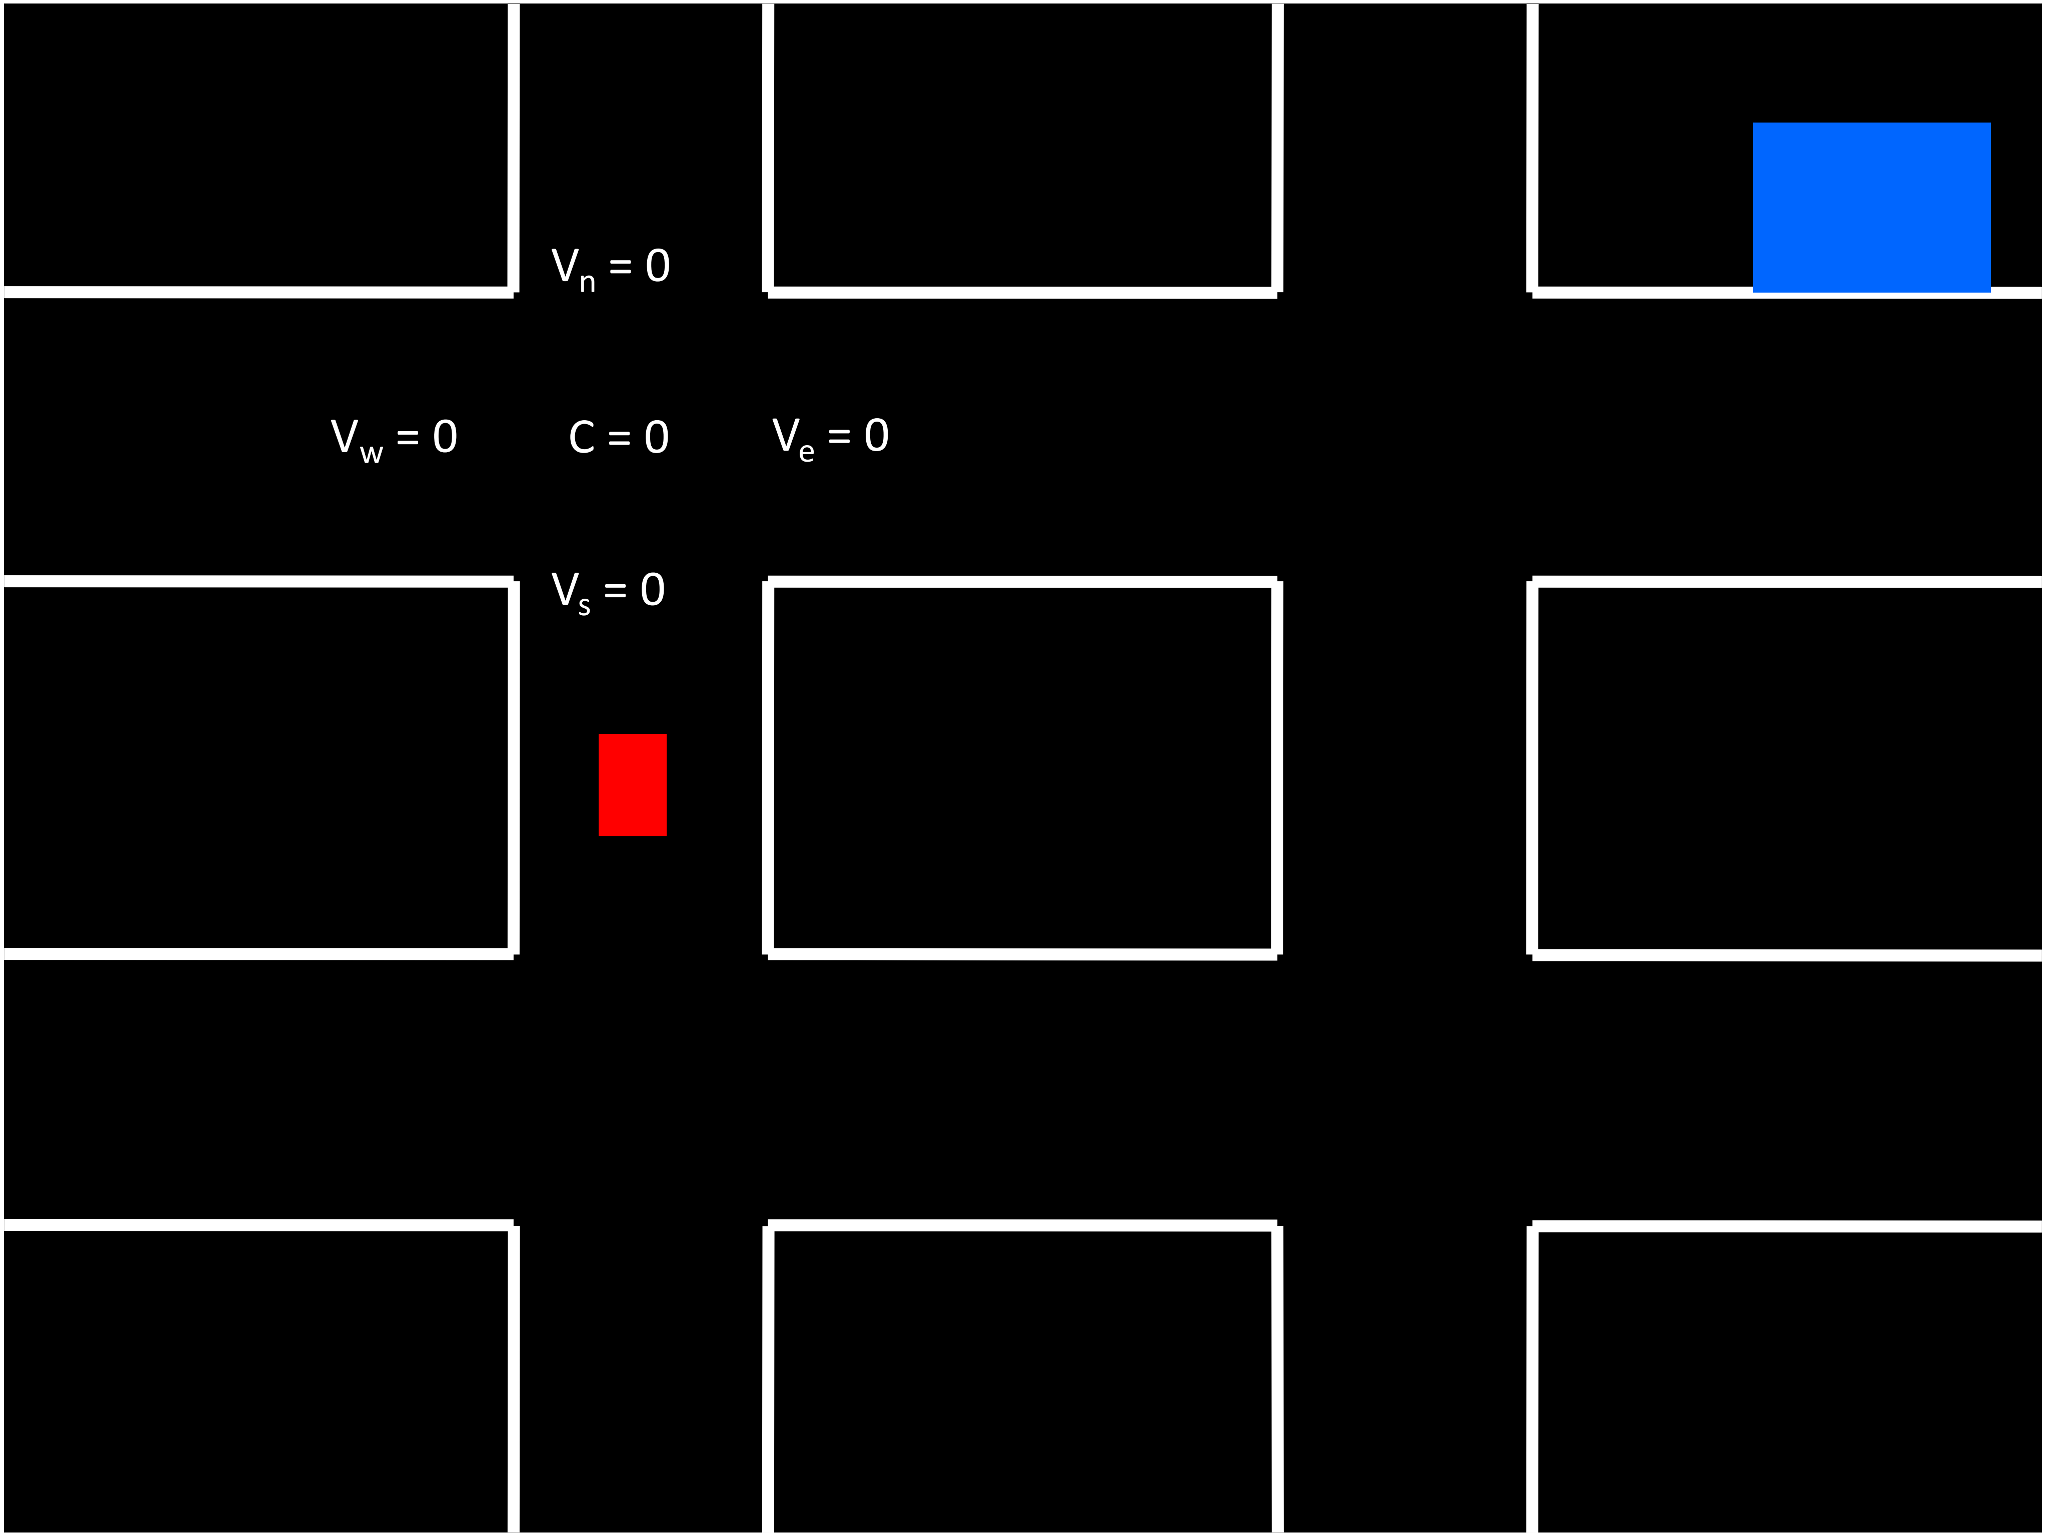

$V_n = 0$
$V_w = 0$
$C = 0$
$V_e = 0$
$V_s = 0$

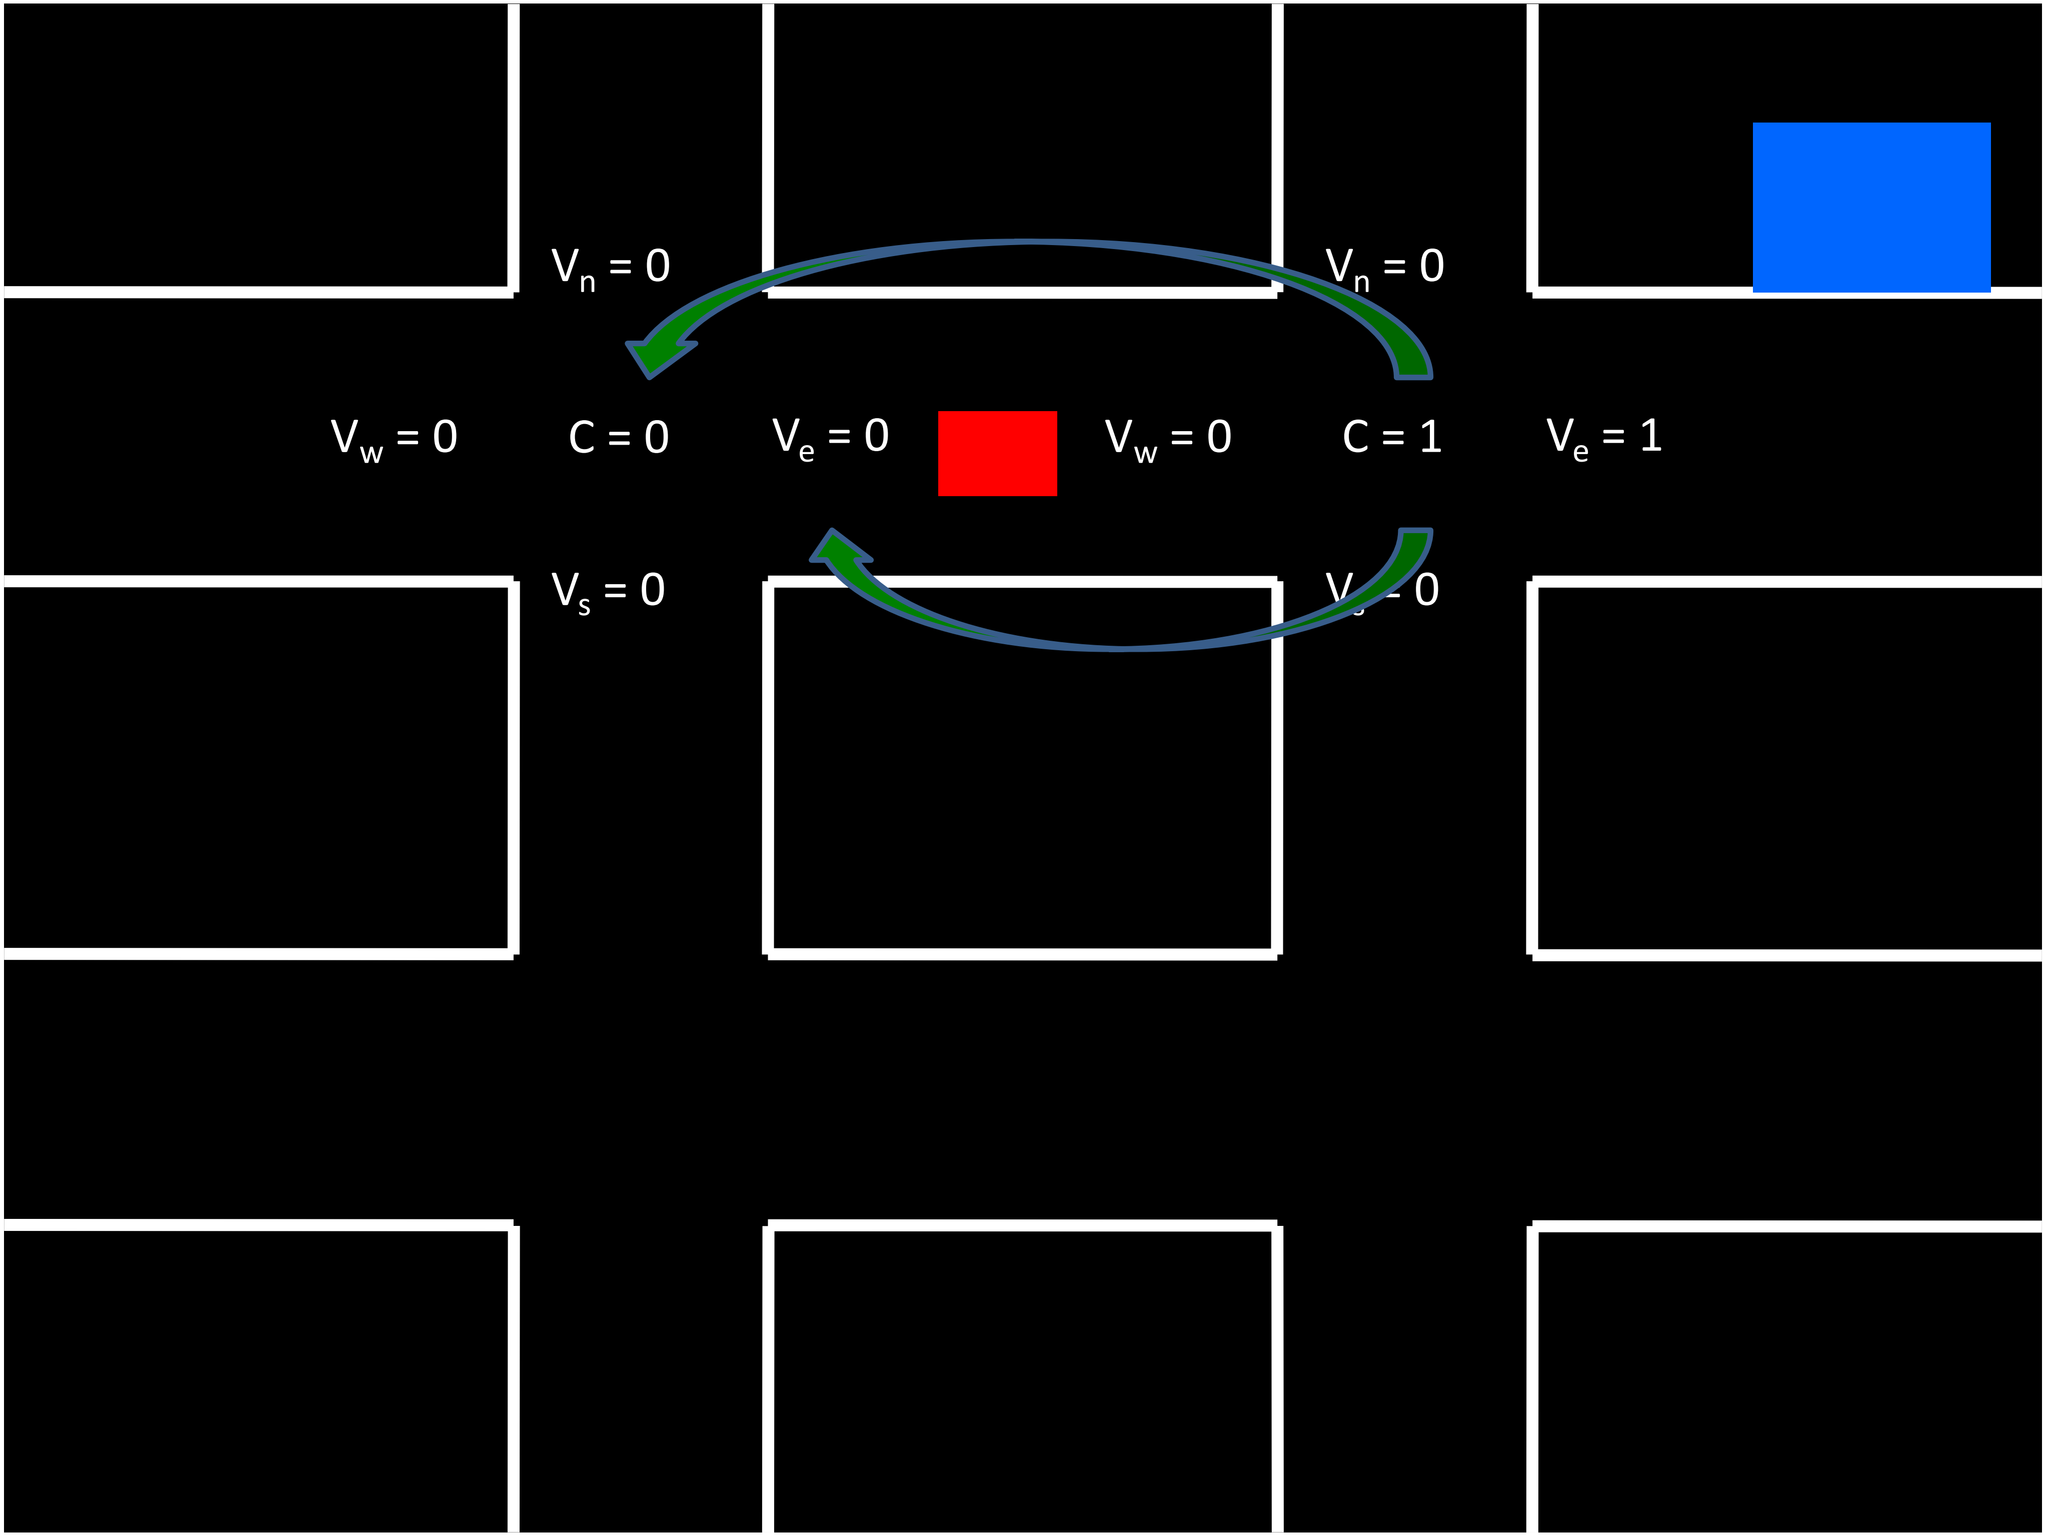

$V_n = 0$
$V_n = 0$
$V_w = 0$
$C = 0$
$V_e = 0$
$V_w = 0$
$C = 1$
$V_e = 1$
$V_s = 0$
$V_s = 0$

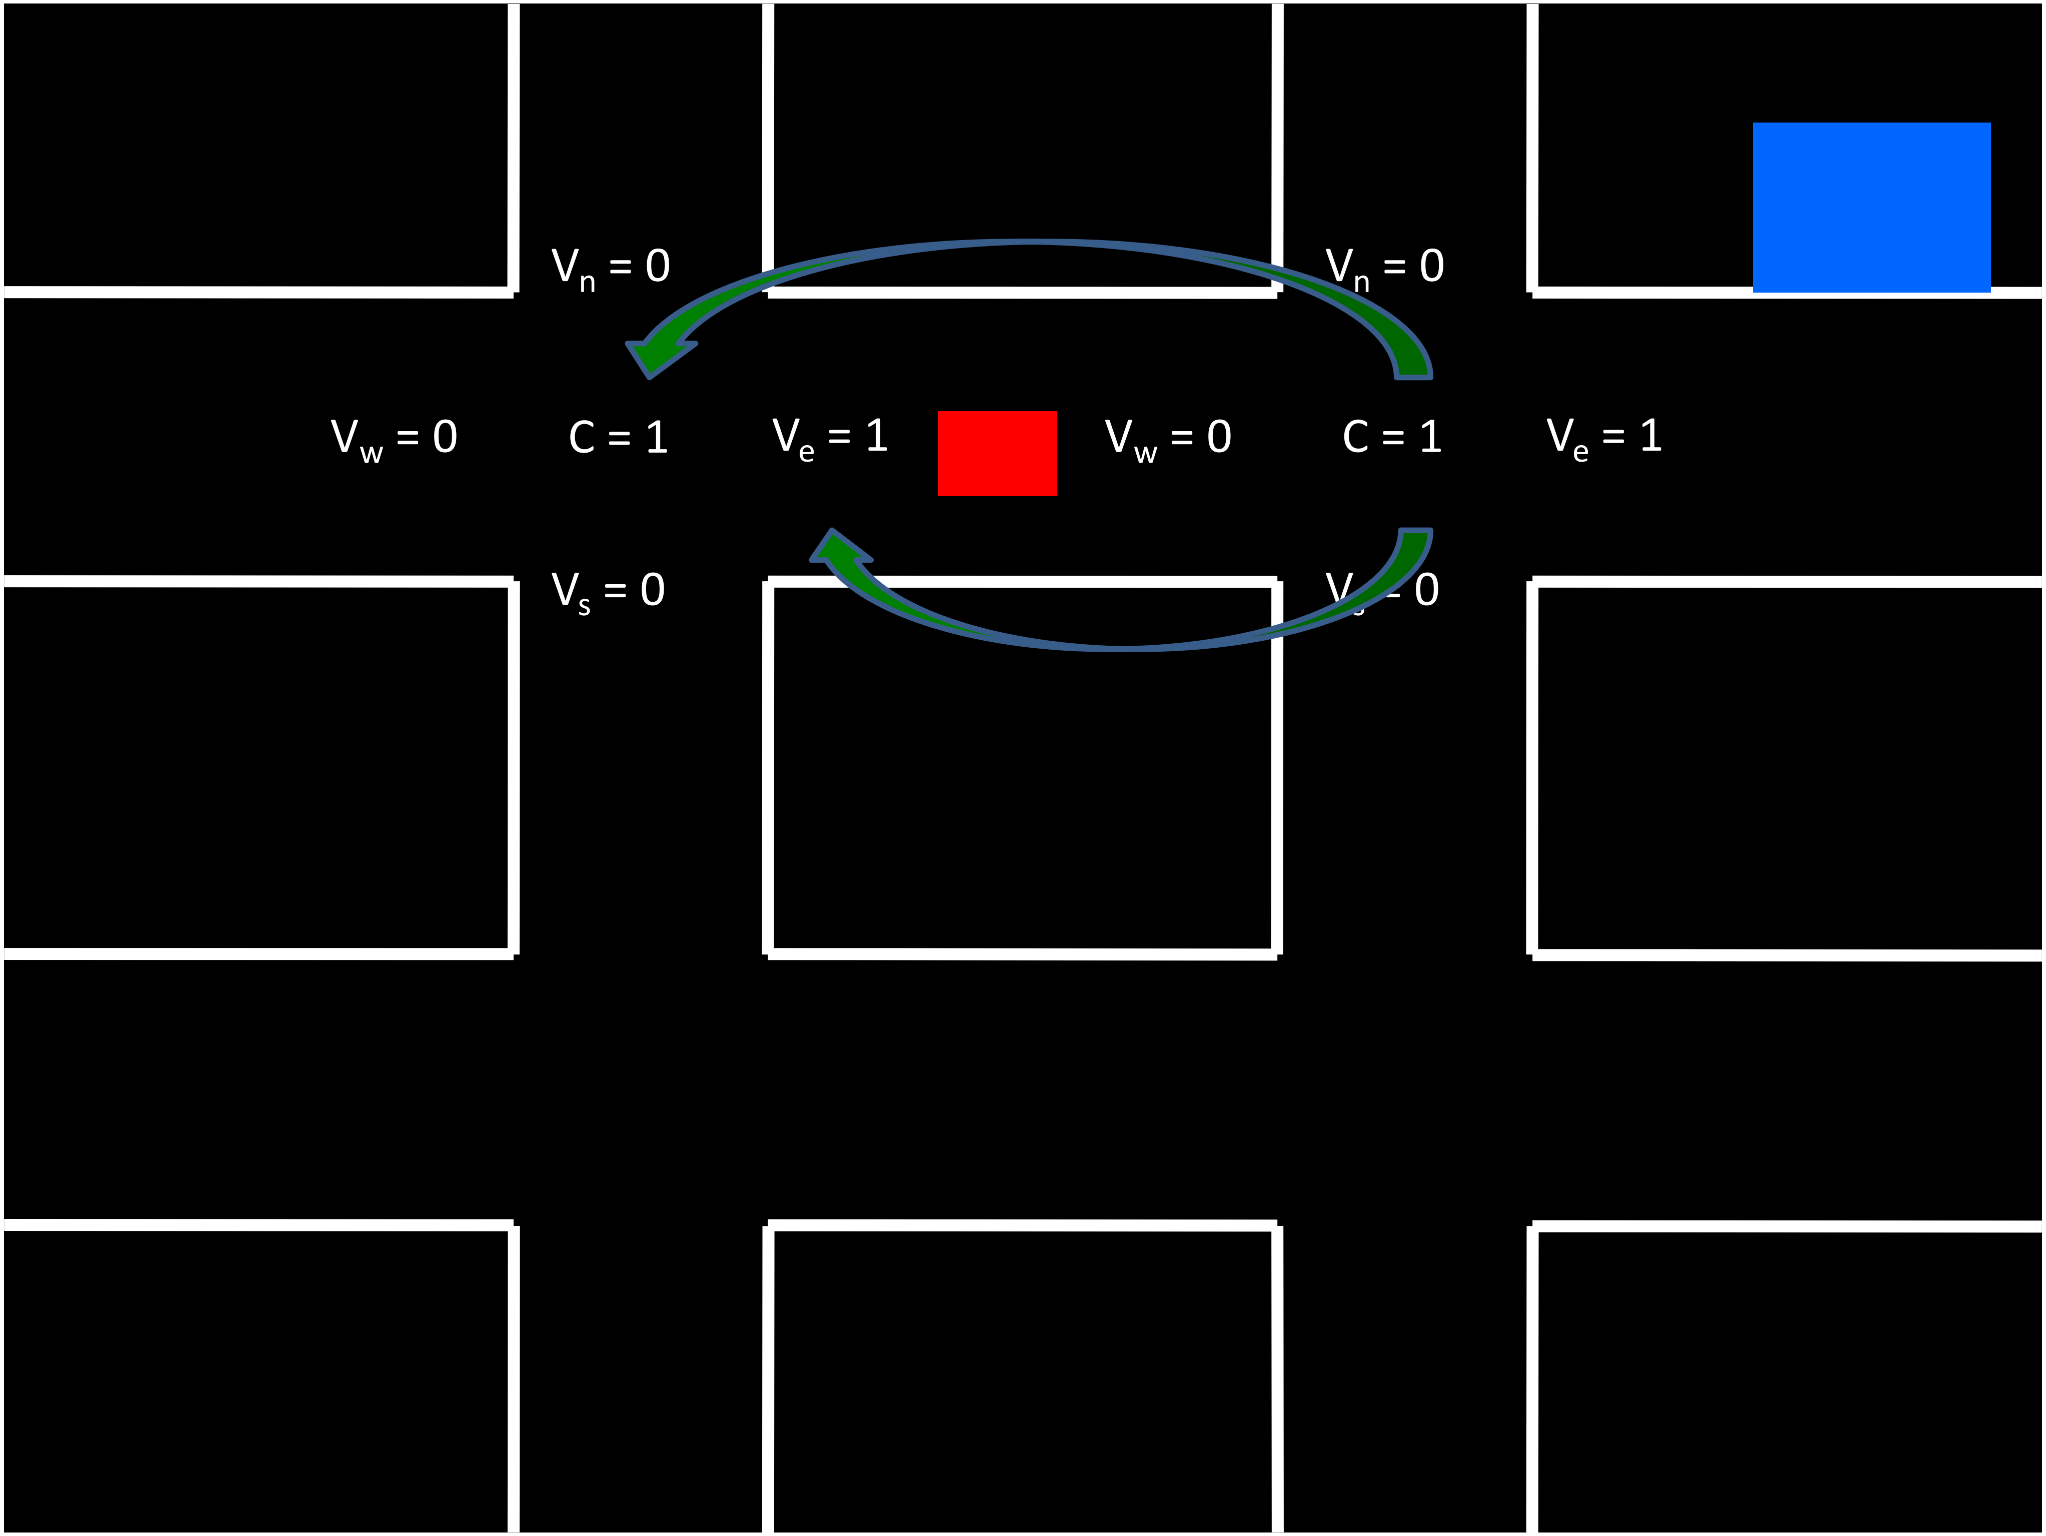

$V_n = 0$
$V_n = 0$
$V_w = 0$
$C = 1$
$V_e = 1$
$V_w = 0$
$C = 1$
$V_e = 1$
$V_s = 0$
$V_s = 0$

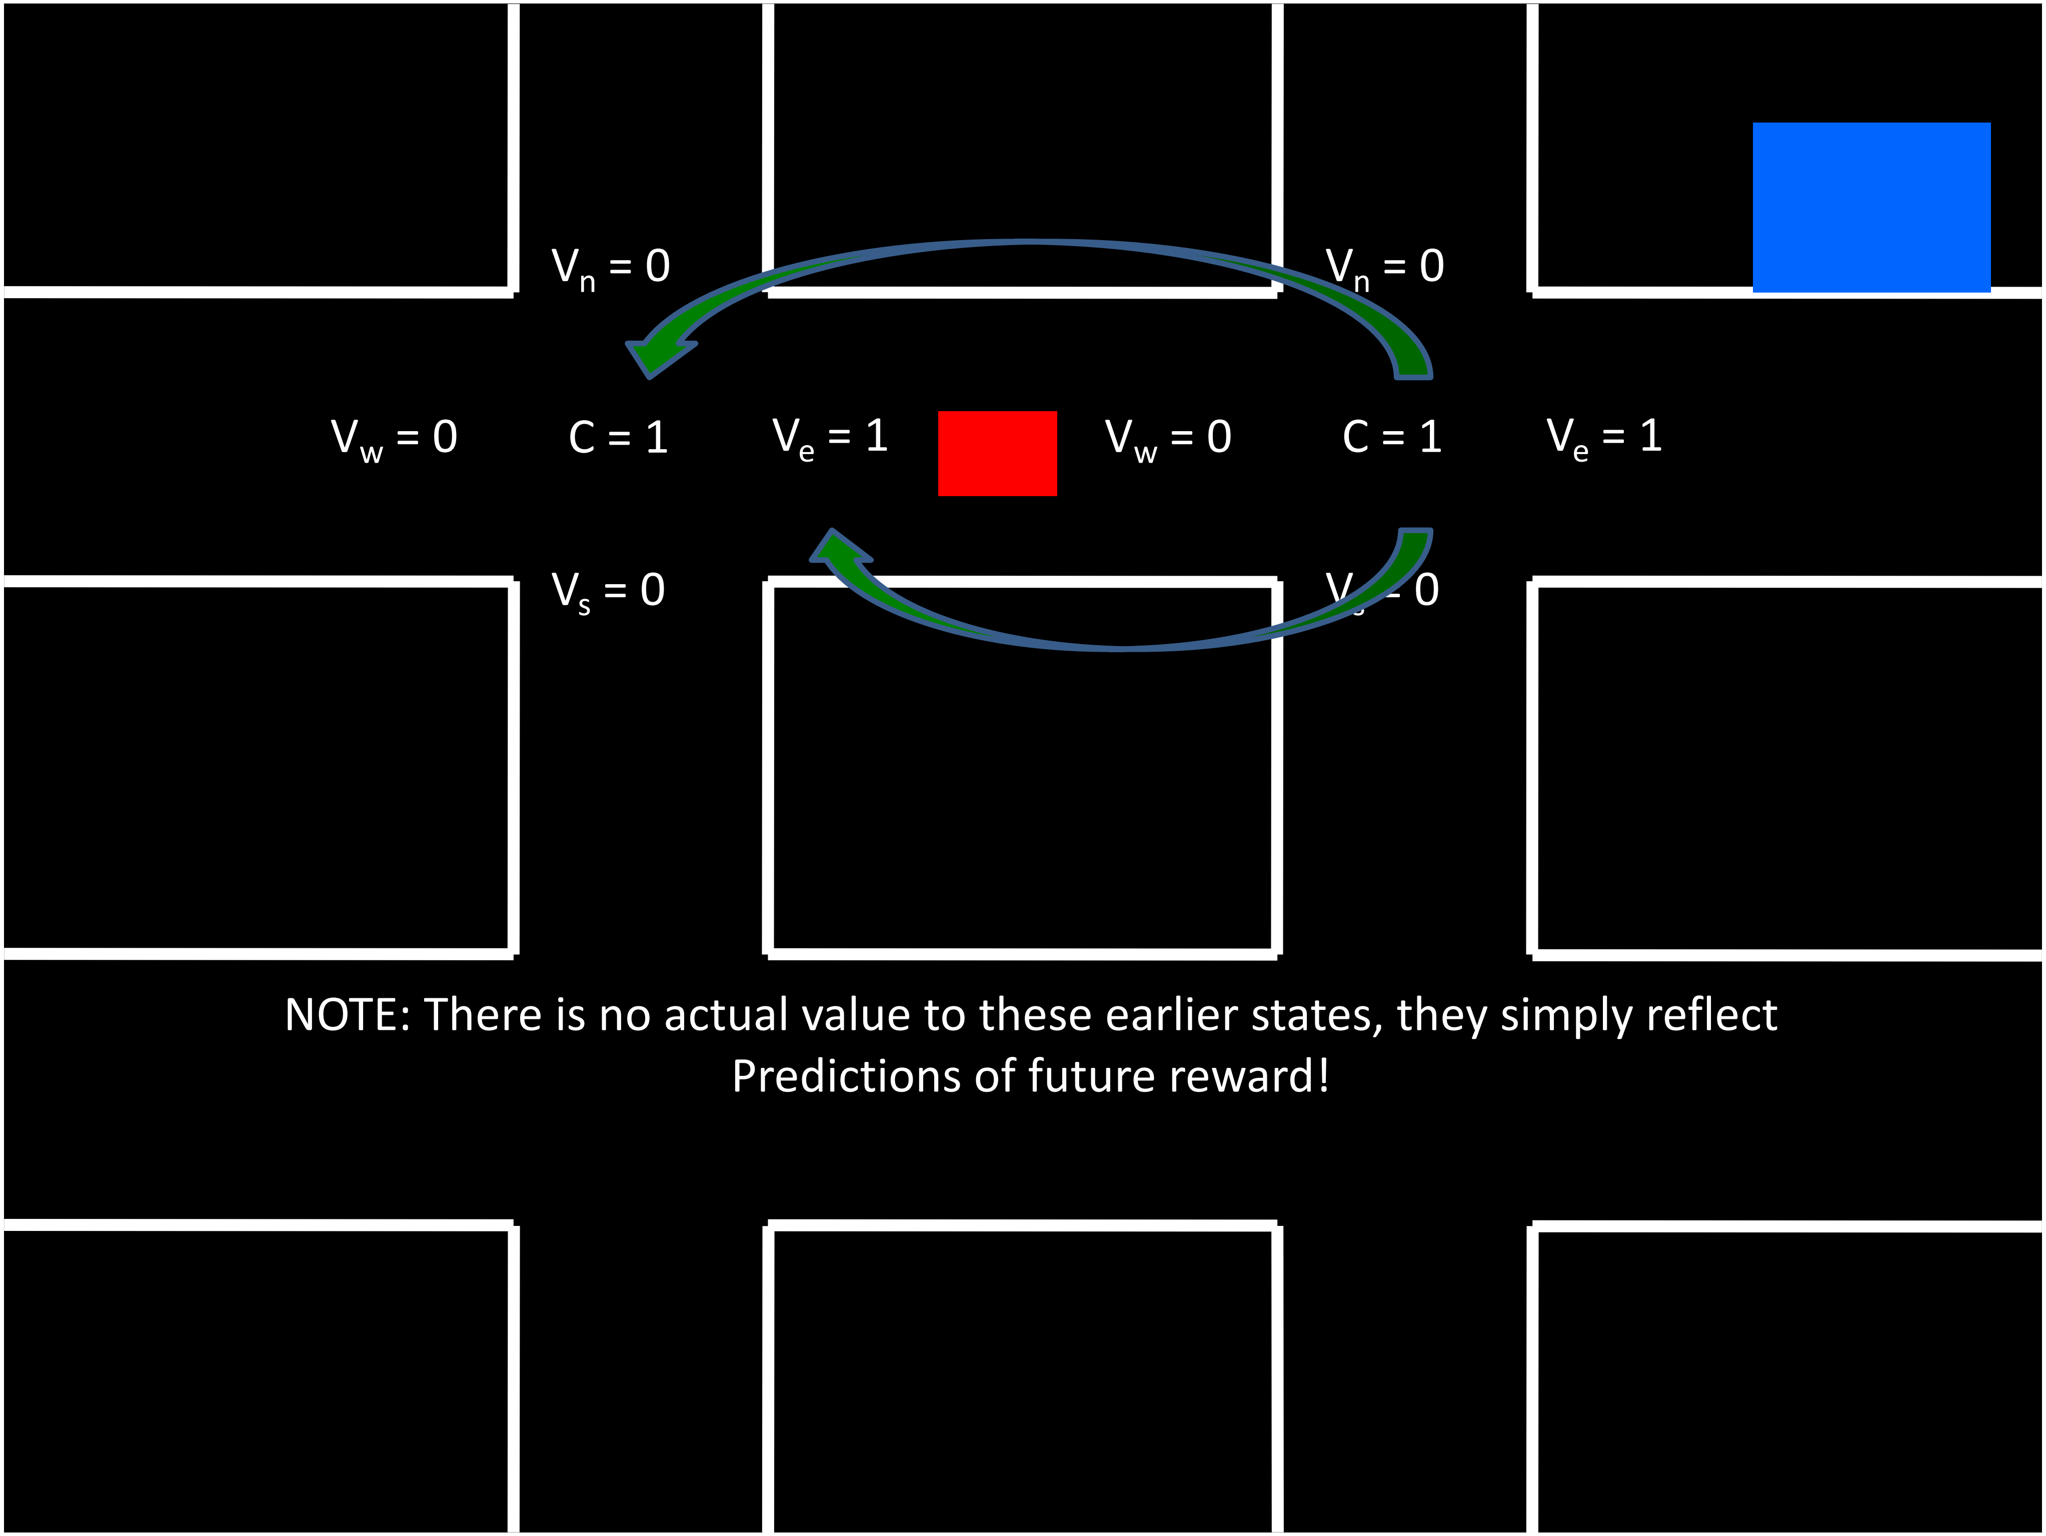

$V_n = 0$
$V_n = 0$
$V_w = 0$
$C = 1$
$V_e = 1$
$V_w = 0$
$C = 1$
$V_e = 1$
$V_s = 0$
$V_s = 0$
NOTE: There is no actual value to these earlier states, they simply reflect
Predictions of future reward!

Tic Tac Toe

X	0.3	0.5
0.3	O	0.1
0.5	0.1	0.8

What do these values really mean?

Tic Tac Toe

X	O	?
O	O	X
O	X	X

Tic Tac Toe

X	O	1.0
O	O	X
O	X	X

But is this enough? Eligibility Traces...

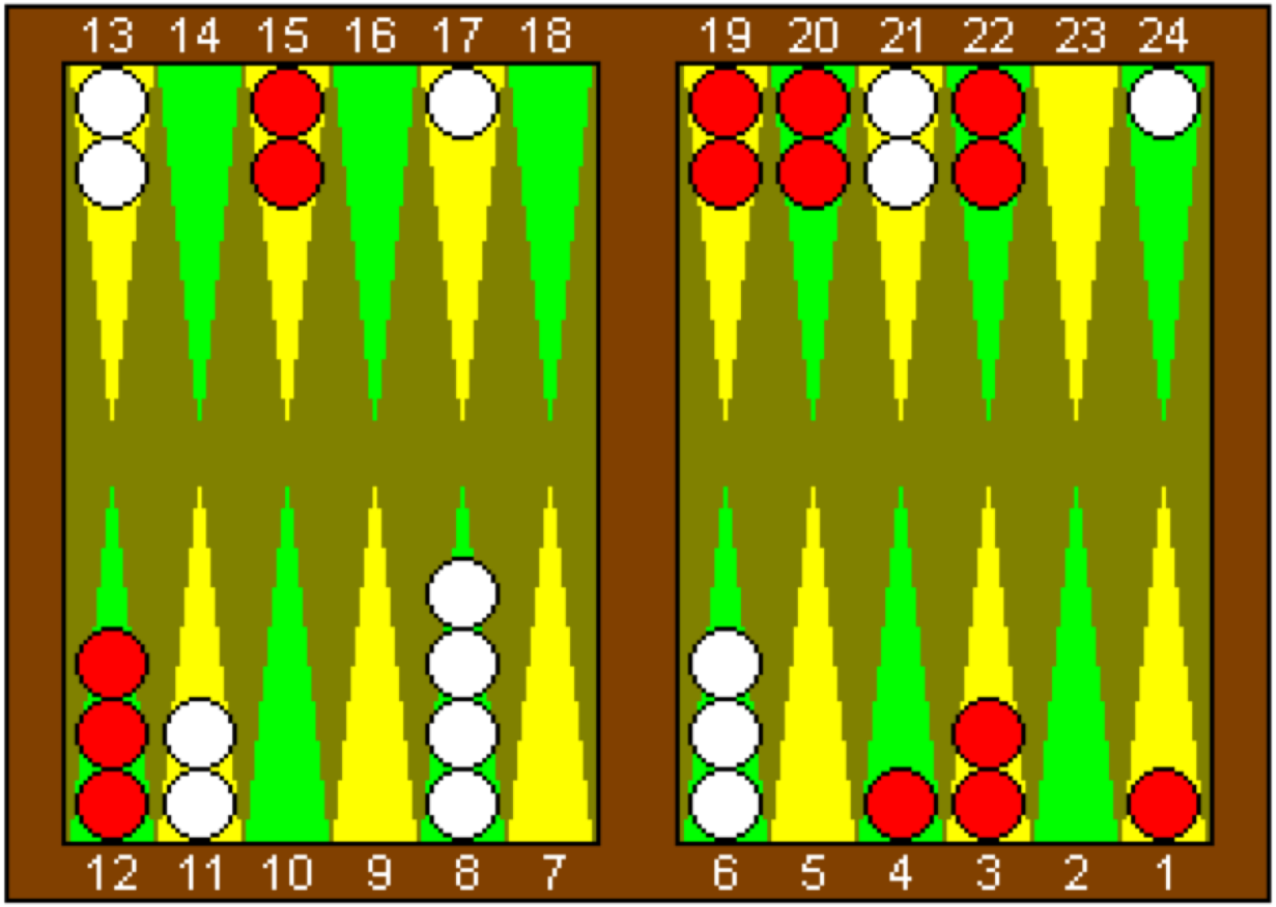

Figure 3. A complex situation where TD-Gammon's positional judgment is apparently superior to traditional expert thinking. White is to play 4-4. The obvious human play is 8-4*, 8-4, 11-7, 11-7. (The asterisk denotes that an opponent checker has been hit.) However, TD-Gammon's choice is the surprising 8-4*, 8-4, 21-17, 21-17! TD-Gammon's analysis of the two plays is given in Table 3.

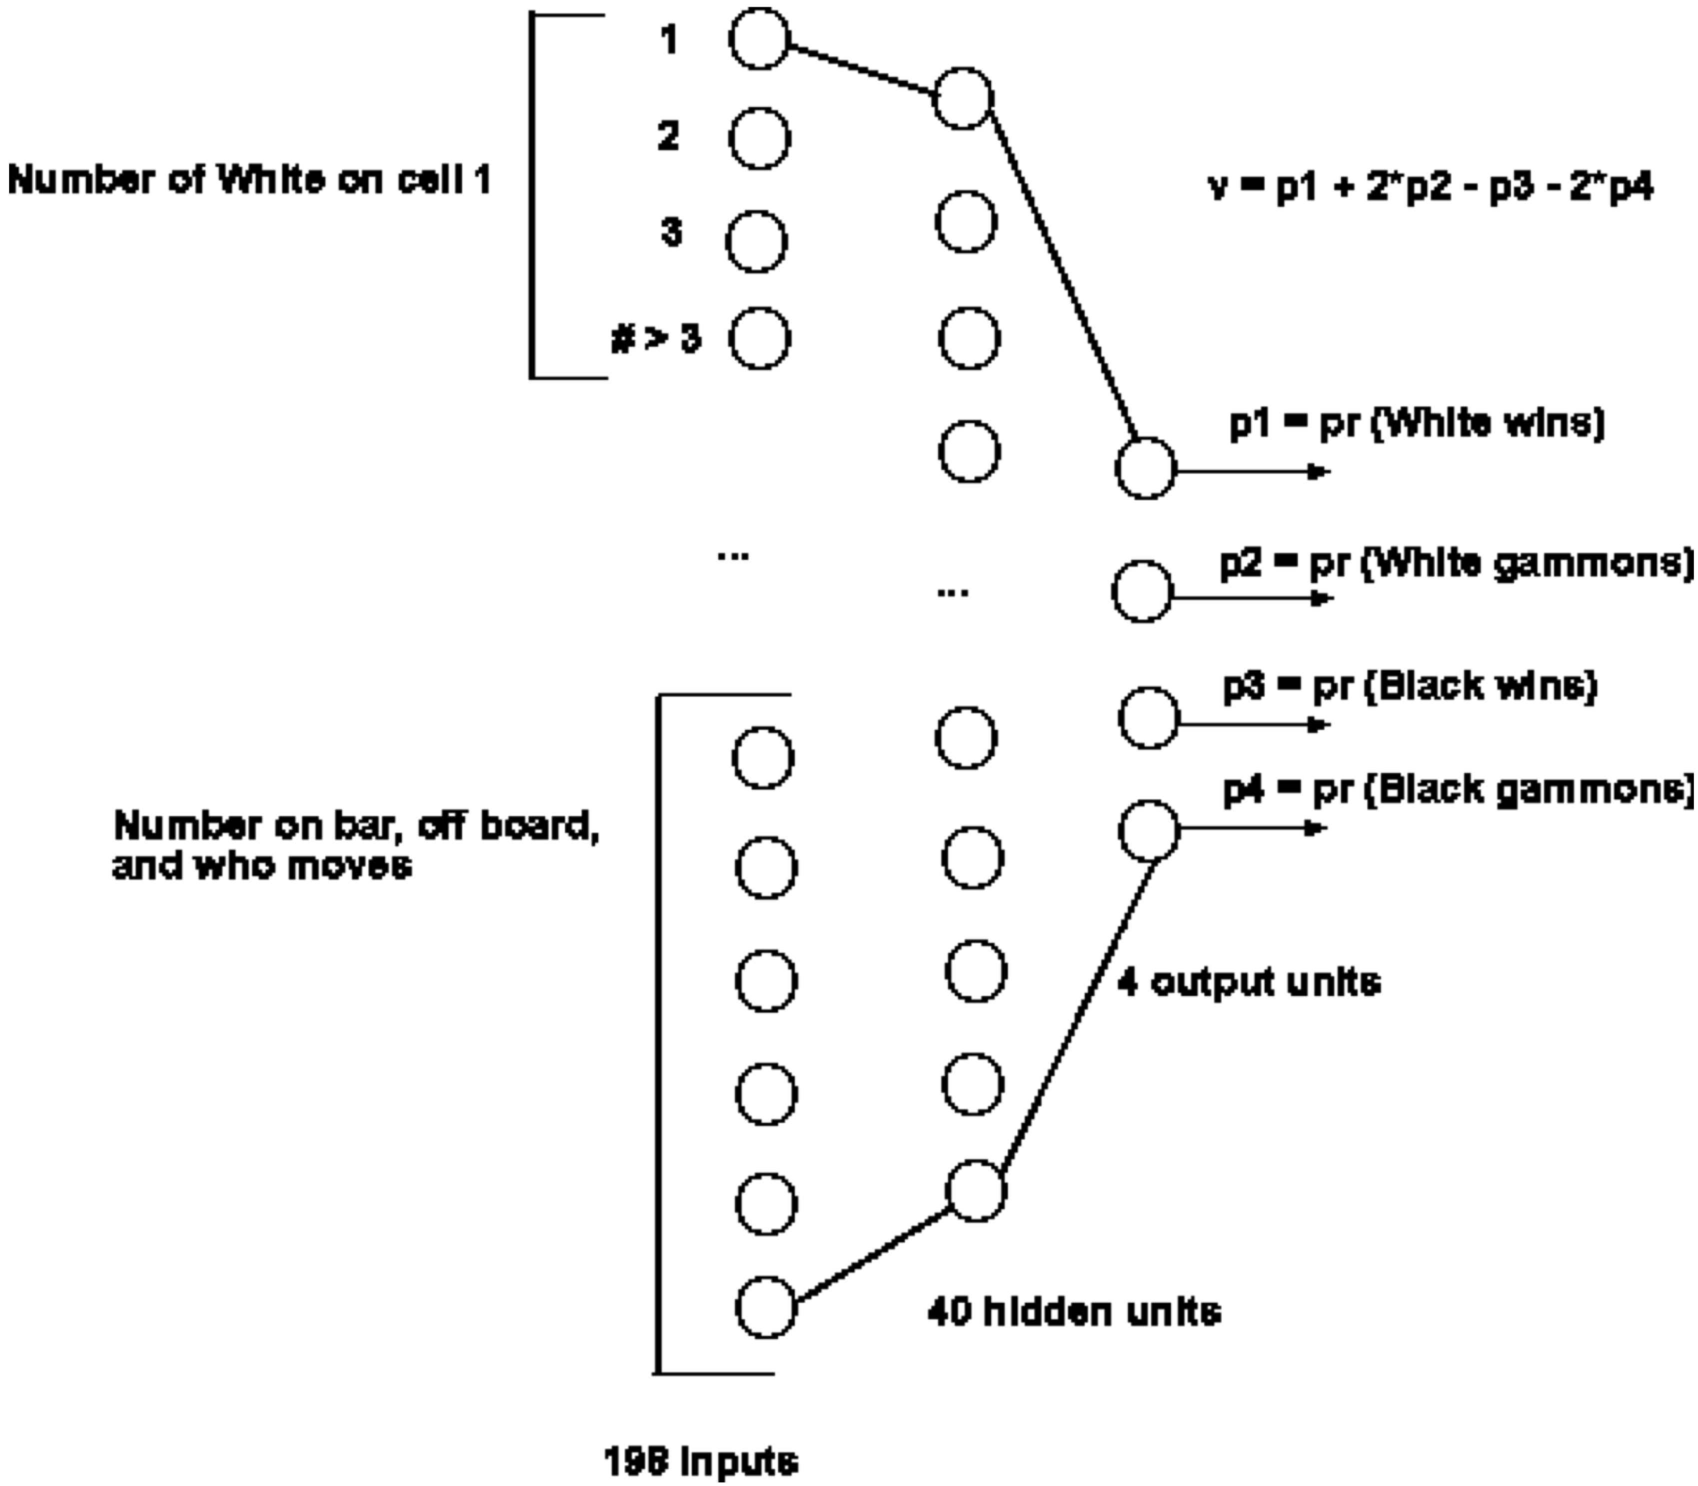

1
2
3
> 3
Number of White on cell 1
v = p1 + 2*p2 - p3 - 2*p4
p1 = pr (White wins)
p2 = pr (White gammons)
p3 = pr (Black wins)
p4 = pr (Black gammons)
...
...
Number on bar, off board, and who moves
4 output units
40 hidden units
198 inputs


STAUNTON

1996 Deep Blue wins a game
1997 Deep Blue wins a match

$p = 0.4345$

p = 0.4511

Modern Methods

1. Lookup Tables
2. Lookup Tables + RL
3. Bayesian Probabilities

REUTERS

Google AI beats Chinese master in ancient game of Go

#SCIENCE NEWS MAY 22, 2017 / 11:54 PM / 6 MONTHS AGO

Google AI beats Chinese master in ancient game of Go

Cate Cadell 4 MIN READ

BEIJING (Reuters) - A Google artificial intelligence program defeated a Chinese grand master at the ancient board game Go on Tuesday, a major feather in the cap for the firm's AI ambitions as it looks to woo Beijing to gain re-entry into the country.

The Problem with Value…

Items for Bid

Matches

Water

Food

Knife

Radio

Tent

Sleeping Bag

Book

Sunglasses

Spending Money

Swimsuit

Suit

Plane Tickets

Rental Car

The Problem with Probability…

What is the best time of day to go to the Doctor’s office?

What is your chance of winning this Black Jack hand?

Rank your chance of dying from…

Car Accident

Bee Sting

Lightning Strike

Plane Crash

Ebola

Terrorist Attack

Shark Attack

1 in 13.3 million	chance of **contracting Ebola in America** this year (based on a model of 12 imported cases of Ebola in the course of a year)
1 in 11 million	chance of **dying in a plane crash** for an American this year
1 in 9.6 million	chance of **dying from a lightning strike** for an American this year
1 in 5.2 million	chance of **dying from a bee sting** for an American this year
1 in 3.7 million	chance of being **killed by a shark** in your lifetime (worldwide)
1 in 9100	chance of being **killed in a car accident** in America this year

The Problem with Huygens

Prospect Theory

Daniel Kahneman and Amos Tversky

- Normative decision-analysis theory assumes that outcomes of decisions (prospects) are described in terms of total wealth
- In reality, however, prospects are considered in terms of gains, losses, or neutral outcomes relative to some *reference point* (the current state)
- In **Prospect Theory** a decision maker considers prospects using a function that values all prospects relative to a reference point.
  - Phase I is framing the decision problem
  - Phase II is evaluating the prospects

Consider…

Problem 1: In addition to whatever you own, you have been given $1000.
You are now asked to choose of these option
50% chance to win $1000 OR get $500 for sure.

Problem 2: In addition to whatever you own, you have been given $2000.
You are now asked to choose of these option
50% change to lose $1000 OR lose $500 for sure.

Kahneman, 2011

Consider…

Problem 1: In addition to whatever you own, you have been given $1000.

You are now asked to choose of these option

50% change to win $1000 OR get $500 for sure

Problem 2: In addition to whatever you own, you have been given $2000.

You are now asked to choose one of these options:

50% chance to lose $1000 OR lose $500 for sure

Kahneman, 2011

1.

Gains

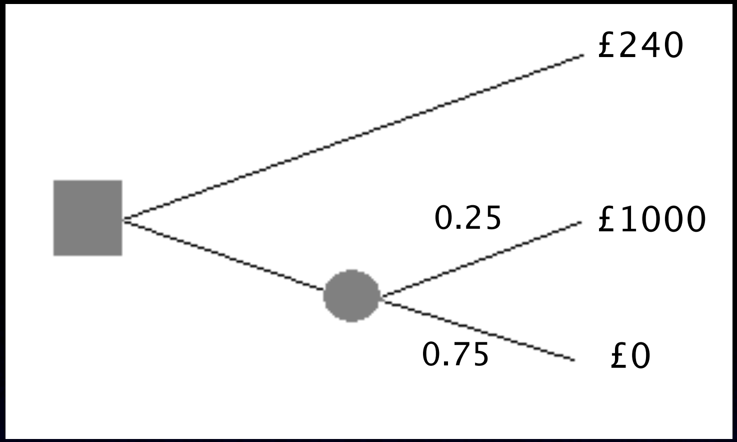

84%

16%

Non-risk taking

2.

Losses

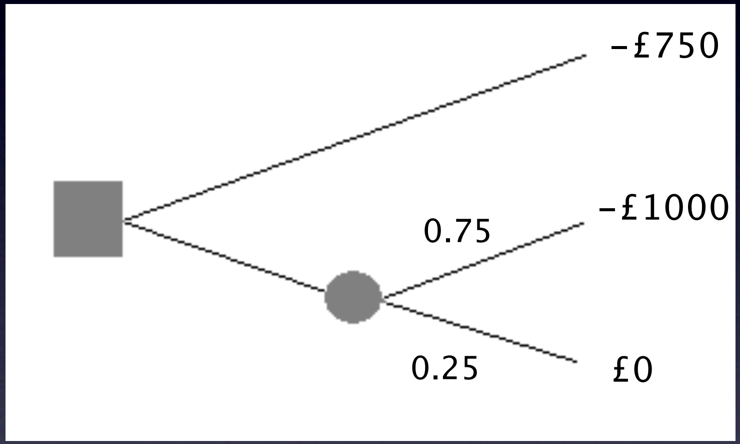

13%

87%

Risk Taking

In the first example individuals are less likely to take the risk in order to win £1000

However when it comes to losses people suddenly become risk takers and more individuals are willing to take the £1000 gamble.

A hypothetical value function

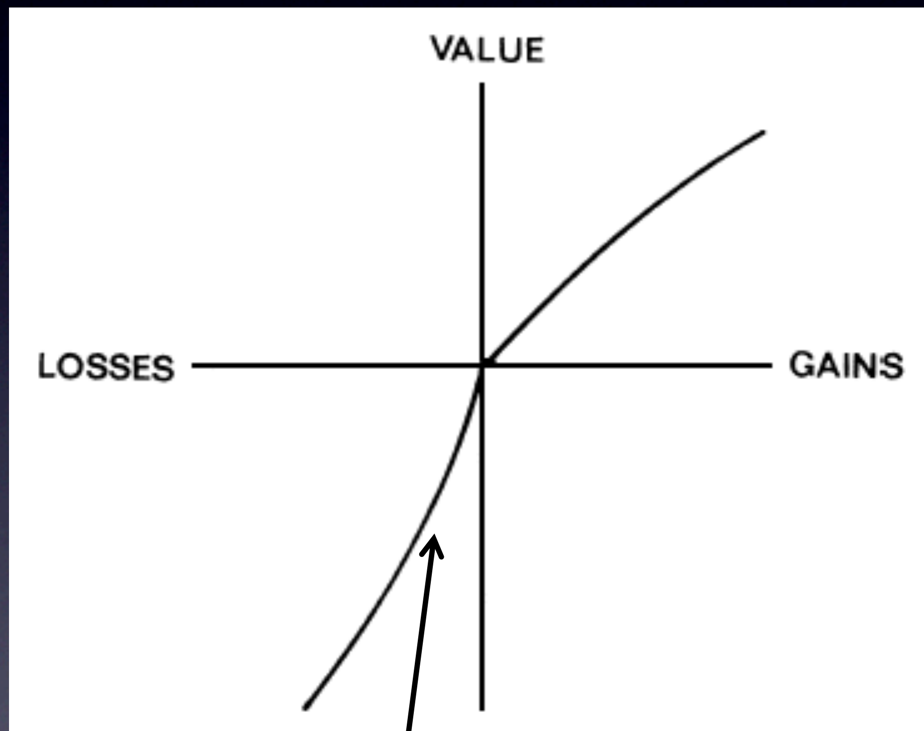

The curve is concave above the line as decision makers will be risk averse when choosing between gains.

The curve convex below the line as decision makers are risk seeking when choosing between losses.

The line at this point is steeper because decision makers are extreme risk takers.

Prospect Theory

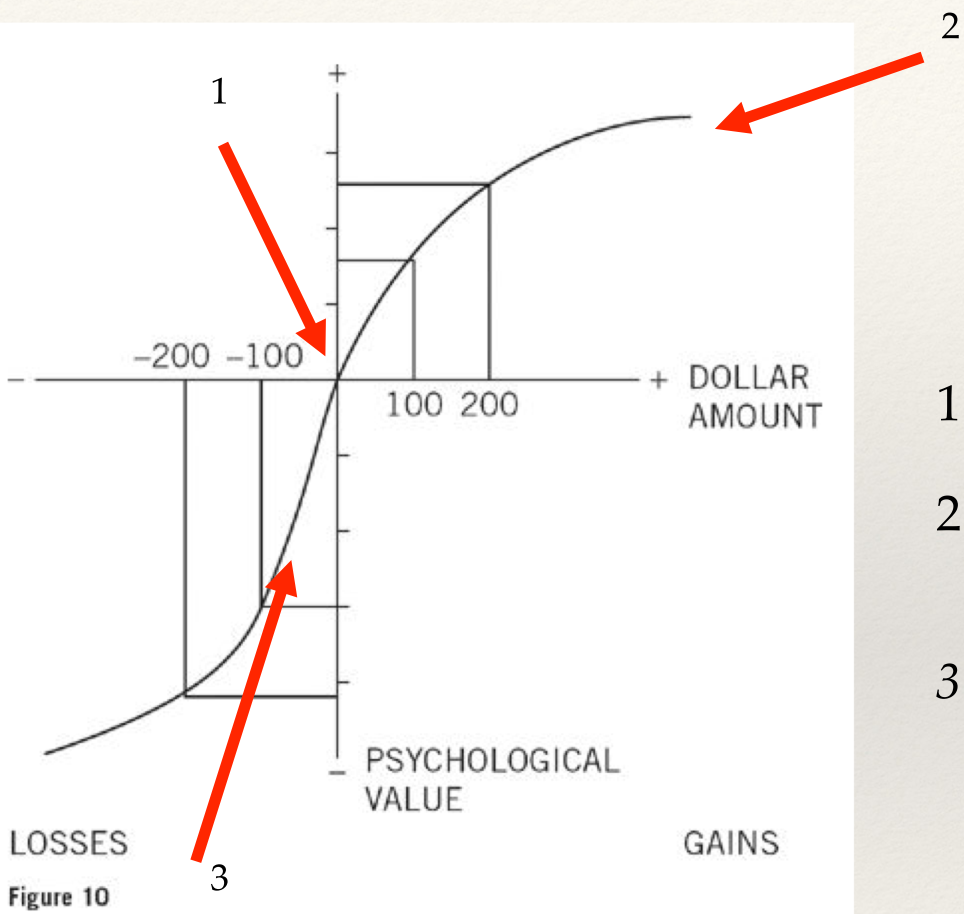

Figure 10

LOSS AVERSION

1. Neutral reference point
2. Diminishing sensitivity to gains and losses
3. *S* is not symmetrical

Prospect Theory: An analysis of decision under risk
Kahneman and Tversky 1979

- A critique of expected utility theory as a descriptive model of decision making under risk, developing an alternative model now known as prospect theory

- The paper presents situations and proof which violate the axioms of expected utility theory

- The theory is based on simple prospects, monetary outcomes and stated probabilities

- But, can be later applied to real life situations

- Evaluates situations involving risk and shows that people respond differently to a risk depending on whether the outcome is a gain or a loss

Properties of Value Functions in Prospect Theory

- The Value function V(*X*), where *X* is a prospect:
  - Is defined by gains and losses from a reference point
  - Is *concave* for gains, and *convex* for losses
    - The value function is steepest near the point of reference: Sensitivity to losses or gains is maximal in the very first unit of gain or loss
  - Is steeper in the *losses* domain than in the *gains* domain
    - Suggests a basic human mechanism (it is easier to make people unhappy than happy)
    - Thus, the negative effect of a *loss* is larger than the positive effect of a *gain*

A Value Function in Prospect Theory

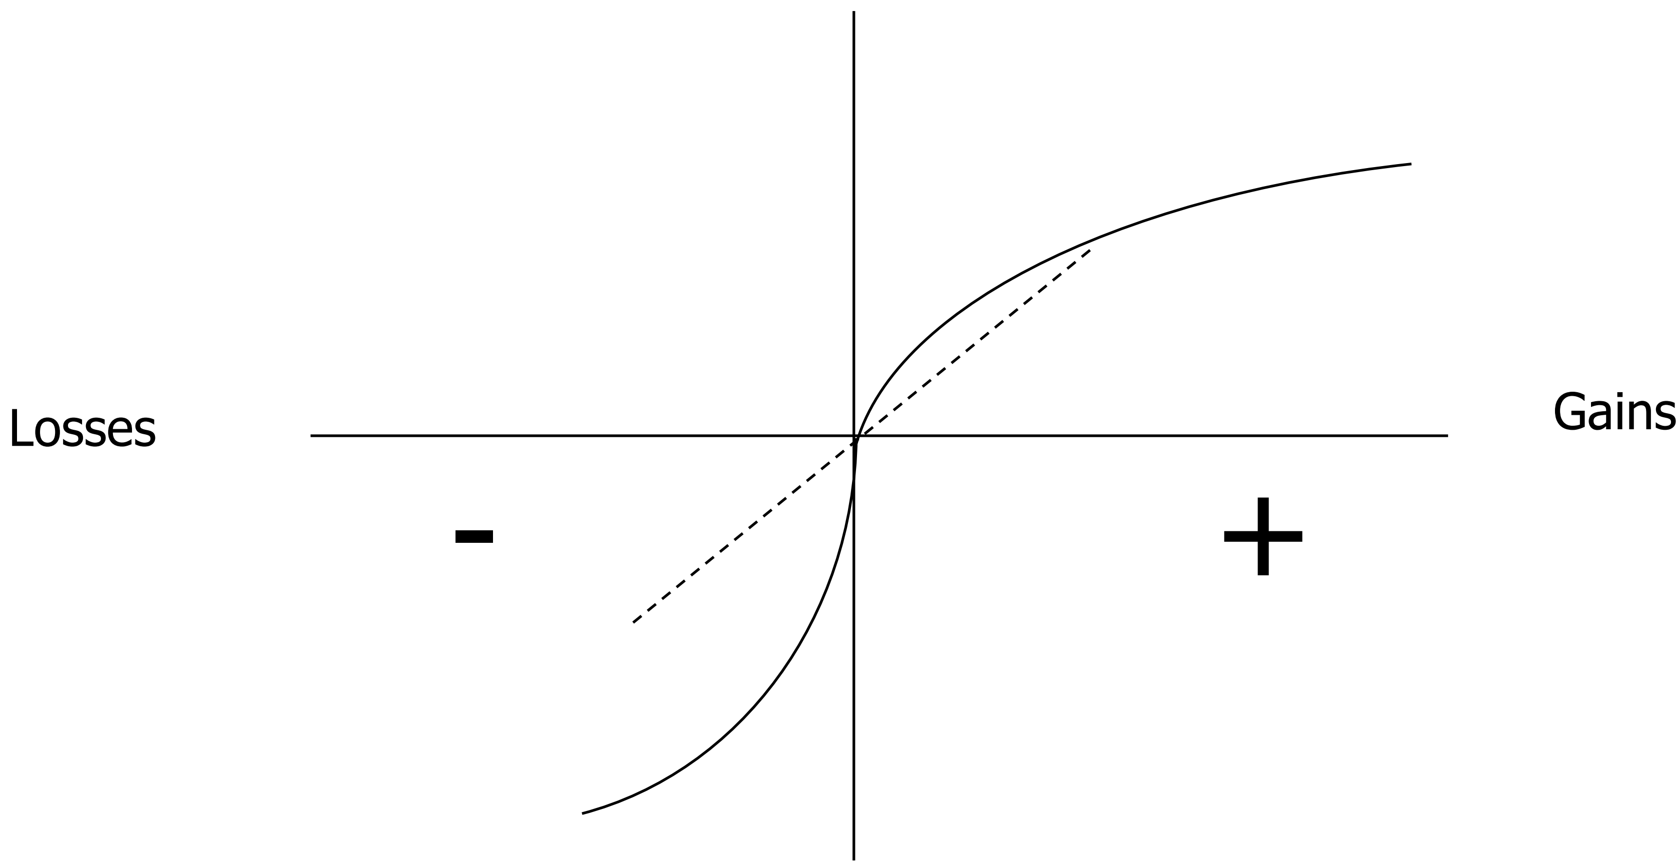

Prospect Theory and Framing

- Risky prospects can be framed in different ways- as gains or as losses
- Changing the description of a prospect should *not* change decisions, but it *does*, in a way predicted by Prospect Theory
- Framing a prospect as a loss rather than a gain, by changing the *reference point*, changes the decision, due to the properties of the value function (which is steeper in the losses domain, i.e., tends to show loss aversion)

Framing: Loss Aversion (I)

- Imagine the US is preparing for the outbreak of an Asian disease, expected to kill 600 people:
  - If program A is adopted, 200 people will be saved (72%)
  - If program B is adopted, there is one third probability that 600 people will be saved and two thirds probability that no people will be saved (28%)

Framing: Loss Aversion (II)

- Imagine the US is preparing for the outbreak of an Asian disease, expected to kill 600 people :
  - If program C is adopted, 400 people will die (28%)
  - If program D is adopted, there is one third probability that nobody will die and two thirds probability that 600 people will die (72%)

Mental Accounts

- Imagine the following situations:
  A. you are about to purchase a jacket for $125 and a calculator for $15. The salesman mentions that the calculator is on sale for $10 at another branch of the store 20 minutes away by car.
  B. you are about to purchase a jacket for $15 and a calculator for $125. The salesman mentions that the calculator is on sale for $120 at another branch of the store 20 minutes away by car.
- 68% of subjects were willing to drive to the other store in A, but only 29% in B

Framing and Mental Accounts

- Three potential framing options, or *accounts*:
  - **Minimal** (considers only differences between local options, such as gaining $5, disregarding common features)
  - **Topical** (considers the context in which the decision arises, such as reducing the price of the calculator from $15 to $10)
  - **Comprehensive** (includes both jacket and calculator in relation to total monthly expenses)
- People usually frame decisions in terms of *topical accounts*; thus the savings on the calculators are considered relative to their prices in each option

Loss Aversion : The Endowment Effect

(Thaler, 1980)

- It is more painful to *give up* an asset than it is pleasurable to *buy* it, an **endowment effect**
  - Thus, *selling* prices are higher than *buying* prices, contrary to economic theory
- There is a ***status quo bias*** (Samuelson and Zeckhauser 1988) since disadvantages of leaving the current state seem larger than advantages

Endowments Effects and Fairness

(Kahneman, Knetch and Thaler, 1986)

- Evaluate the two situations for fairness:
  - A shortage had developed for a car model and the dealer now prices it $200 above list price
  - A shortage had developed for a car model and the dealer, who had been giving $200 discounts, now sells the model for list price
- Perception of fairness depends on whether the question is framed as a reduction in gain or as an actual loss
- *Imposing a surcharge* is considered less fair than *eliminating a discount*; thus, cash prices are presented as a *discount*, rather than credit prices a *surcharge*

Framing: Losses Versus Costs

- Consider the following two propositions:
  A. A gamble that offers a 10% chance of winning \$95 and a 90% chance of losing \$5
  B. A payment of \$5 to participate in a lottery with a 10% chance of winning \$100 and a 90% chance of winning nothing
- 55 of 132 subjects reversed their preferences, 42 of them rejecting A and accepting B
- Thinking of \$5 as payment (cost) rather than a loss makes the venture more acceptable
- Similarly in experiments in which insurance is presented as the cost of protection vs. a sure loss

Framing: Sunk-Cost (Dead Loss) Effect

- Framing negative outcomes as costs rather than as a loss improves subjective feelings, as in:
  - Playing with a tennis elbow in a tennis club, in spite of pain, maintains evaluation of the membership fee as a cost rather than a dead loss
  - Continuing a project that has already cost a lot without any results, rather than starting a new one, although the previous costs are sunk costs
  - Eating bad-tasting food you have already paid for

Summary

- Understanding human behavior requires
  - separation of prospects into positive and negative, relative to a status quo, rather than considering absolute values, and
- Loss aversion, endowment effects, and status quo bias predict (correctly) less symmetry and reversibility in the world than do normative theories
- Framing changes the way decision makers treat the same problem, as can be predicted when the nature of these "anomalies" is taken into account

Another Issue

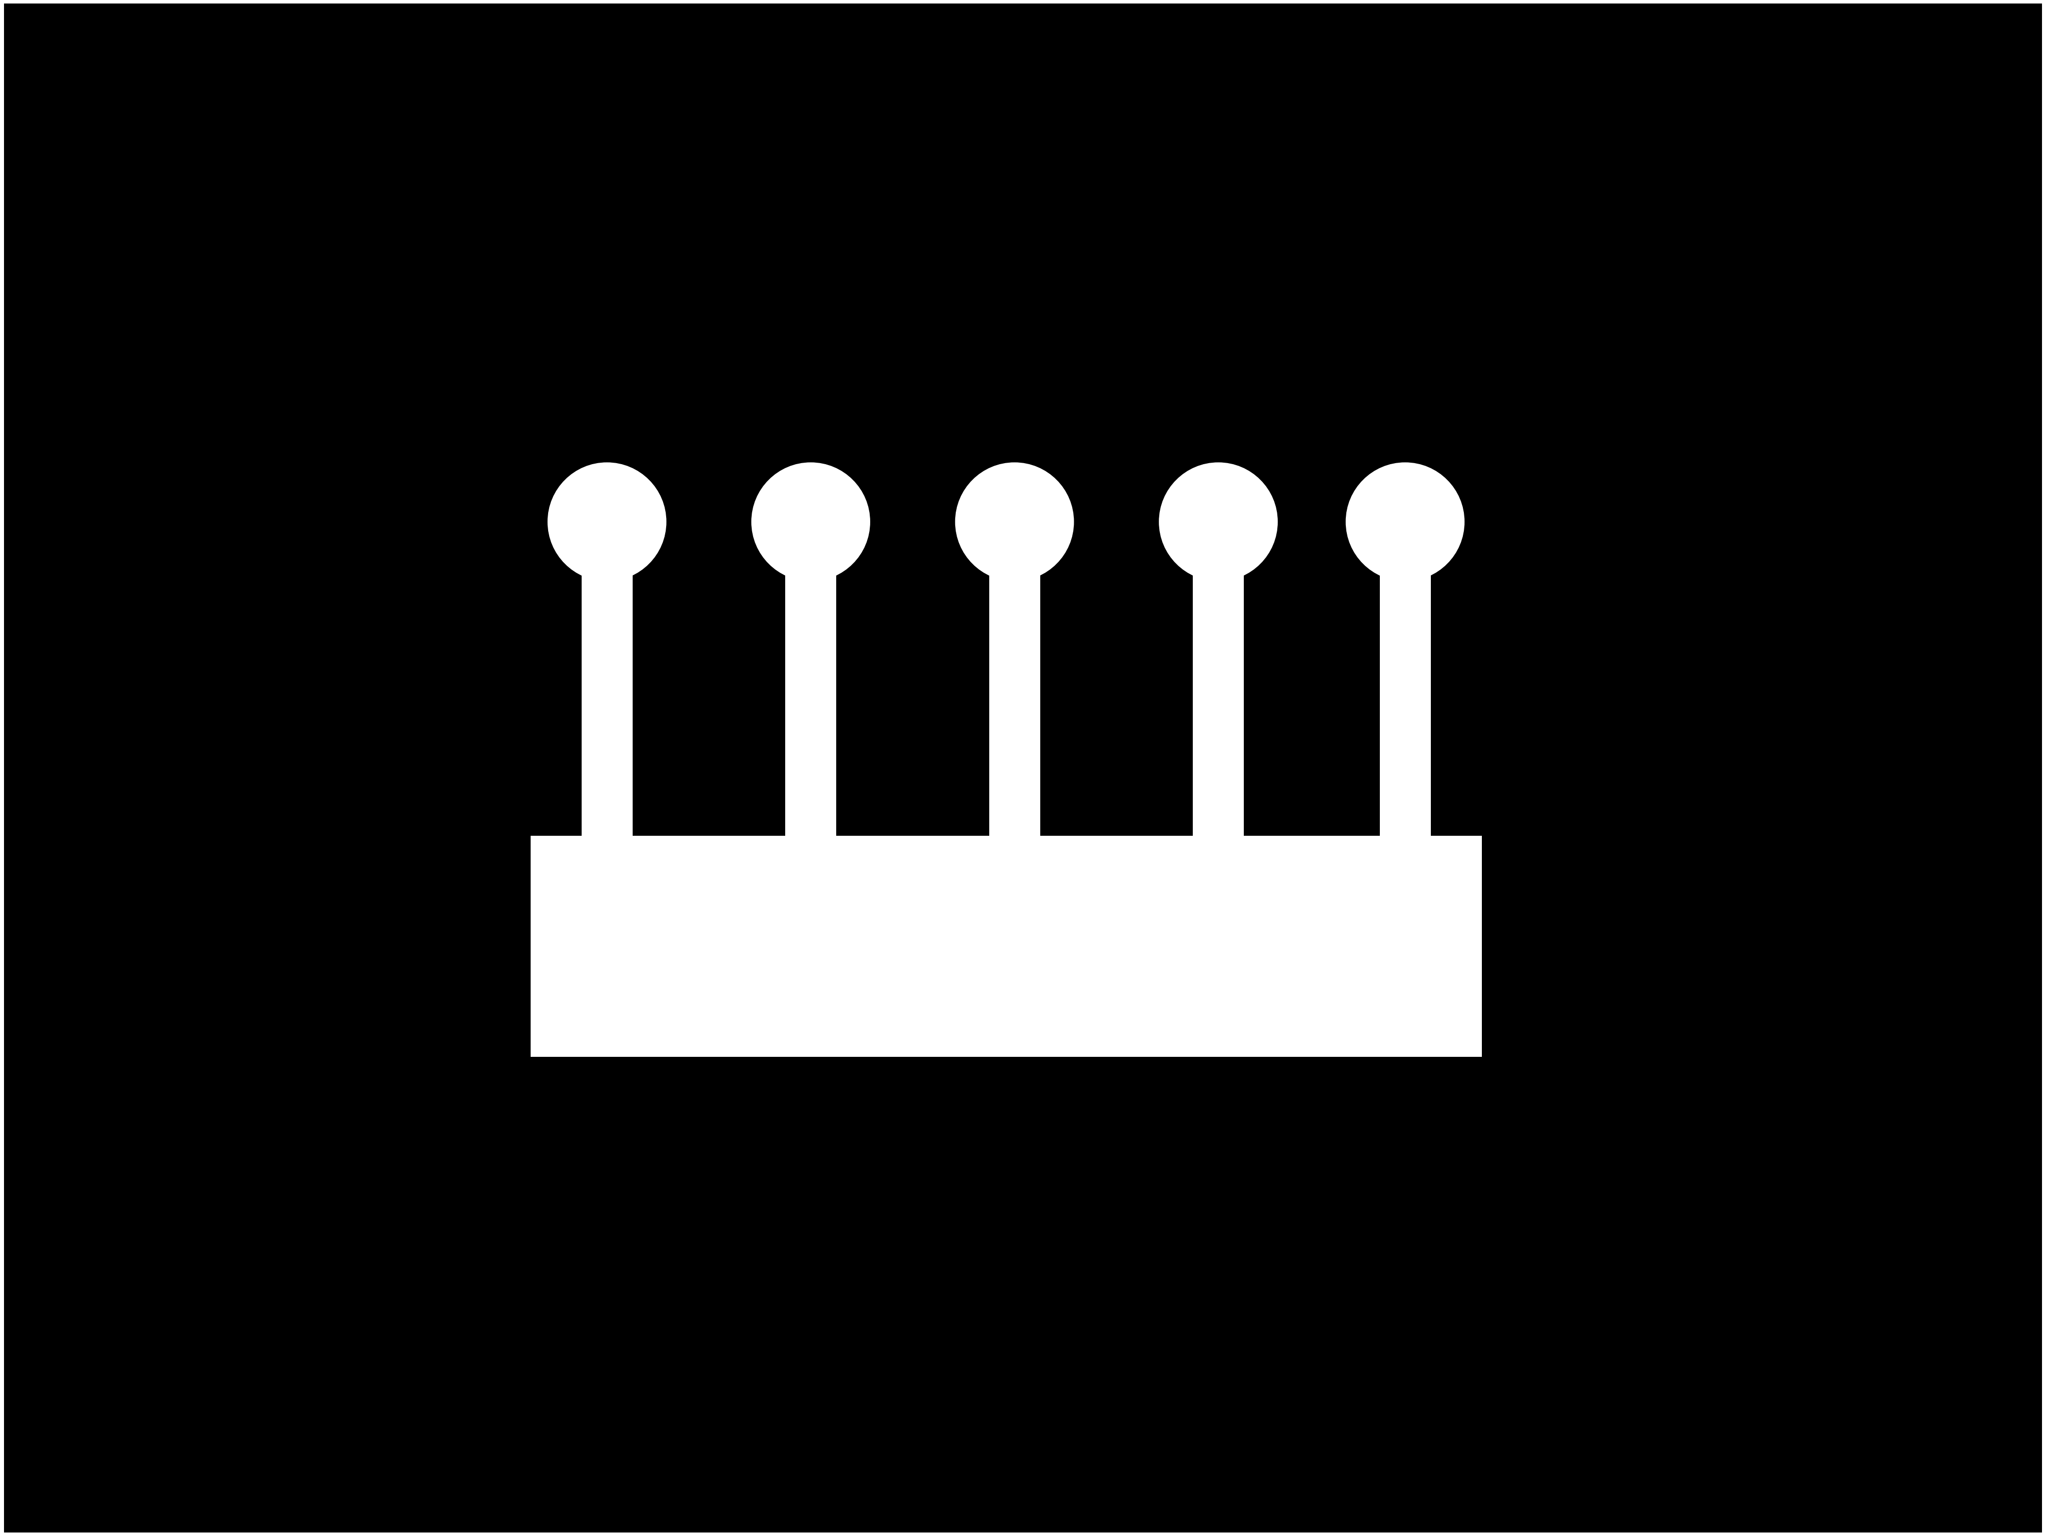

$

?

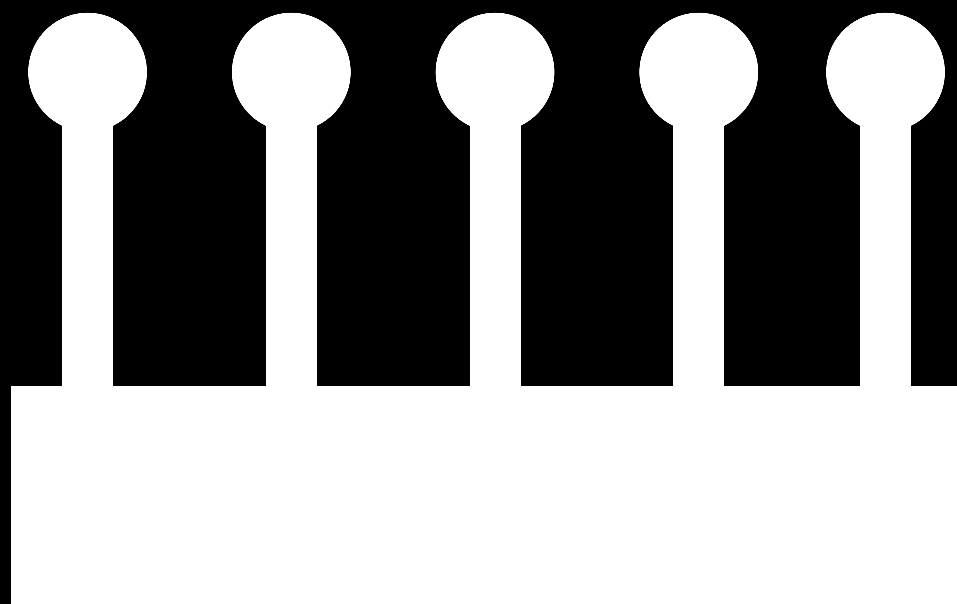

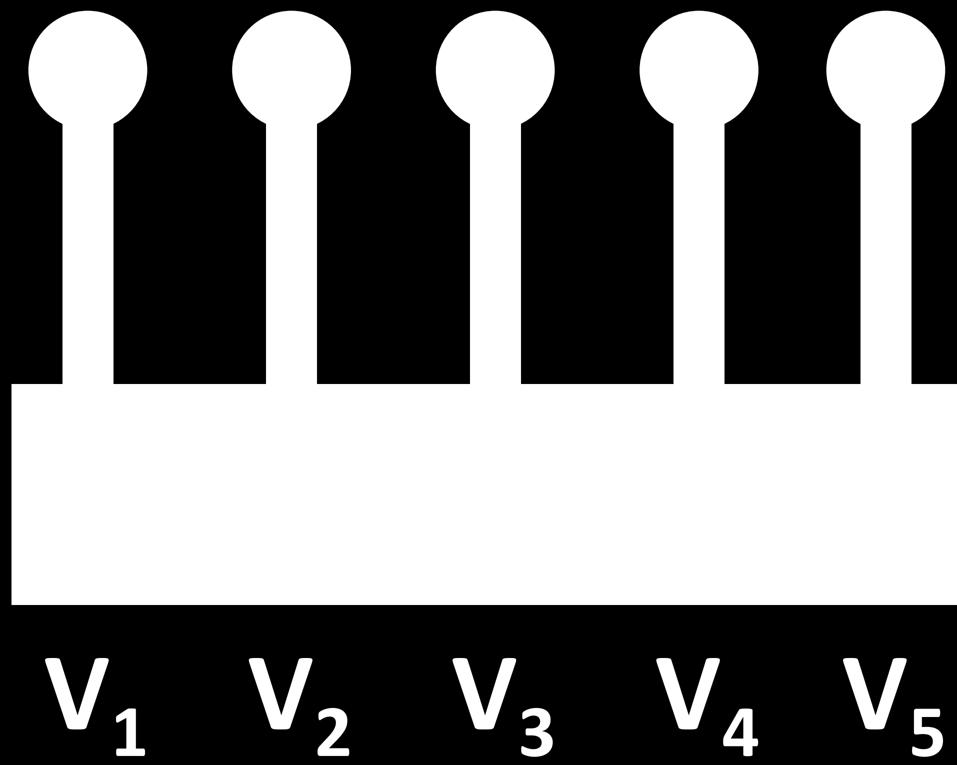

Where: $V_n = \sum r_n / \text{attempts}_n$

Decision Making

1. Always choose the highest value option

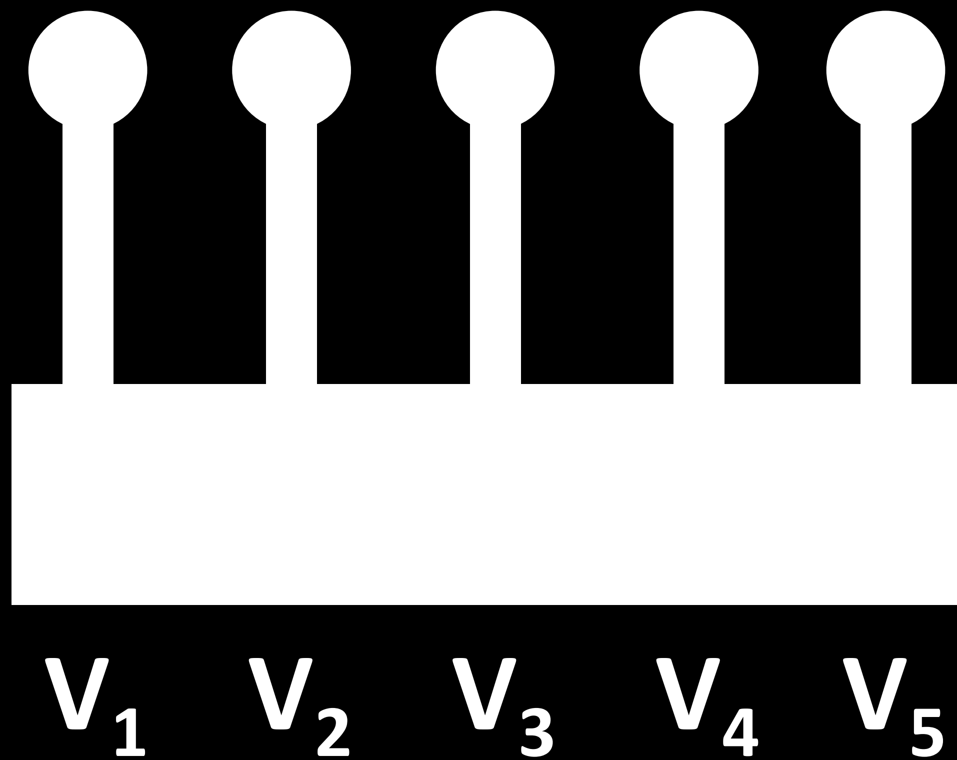

Where: $V_n = \sum r_n / attempts_n$

Decision Making

1. Always Choose the Highest Value Option
2. Exploration versus Exploitation

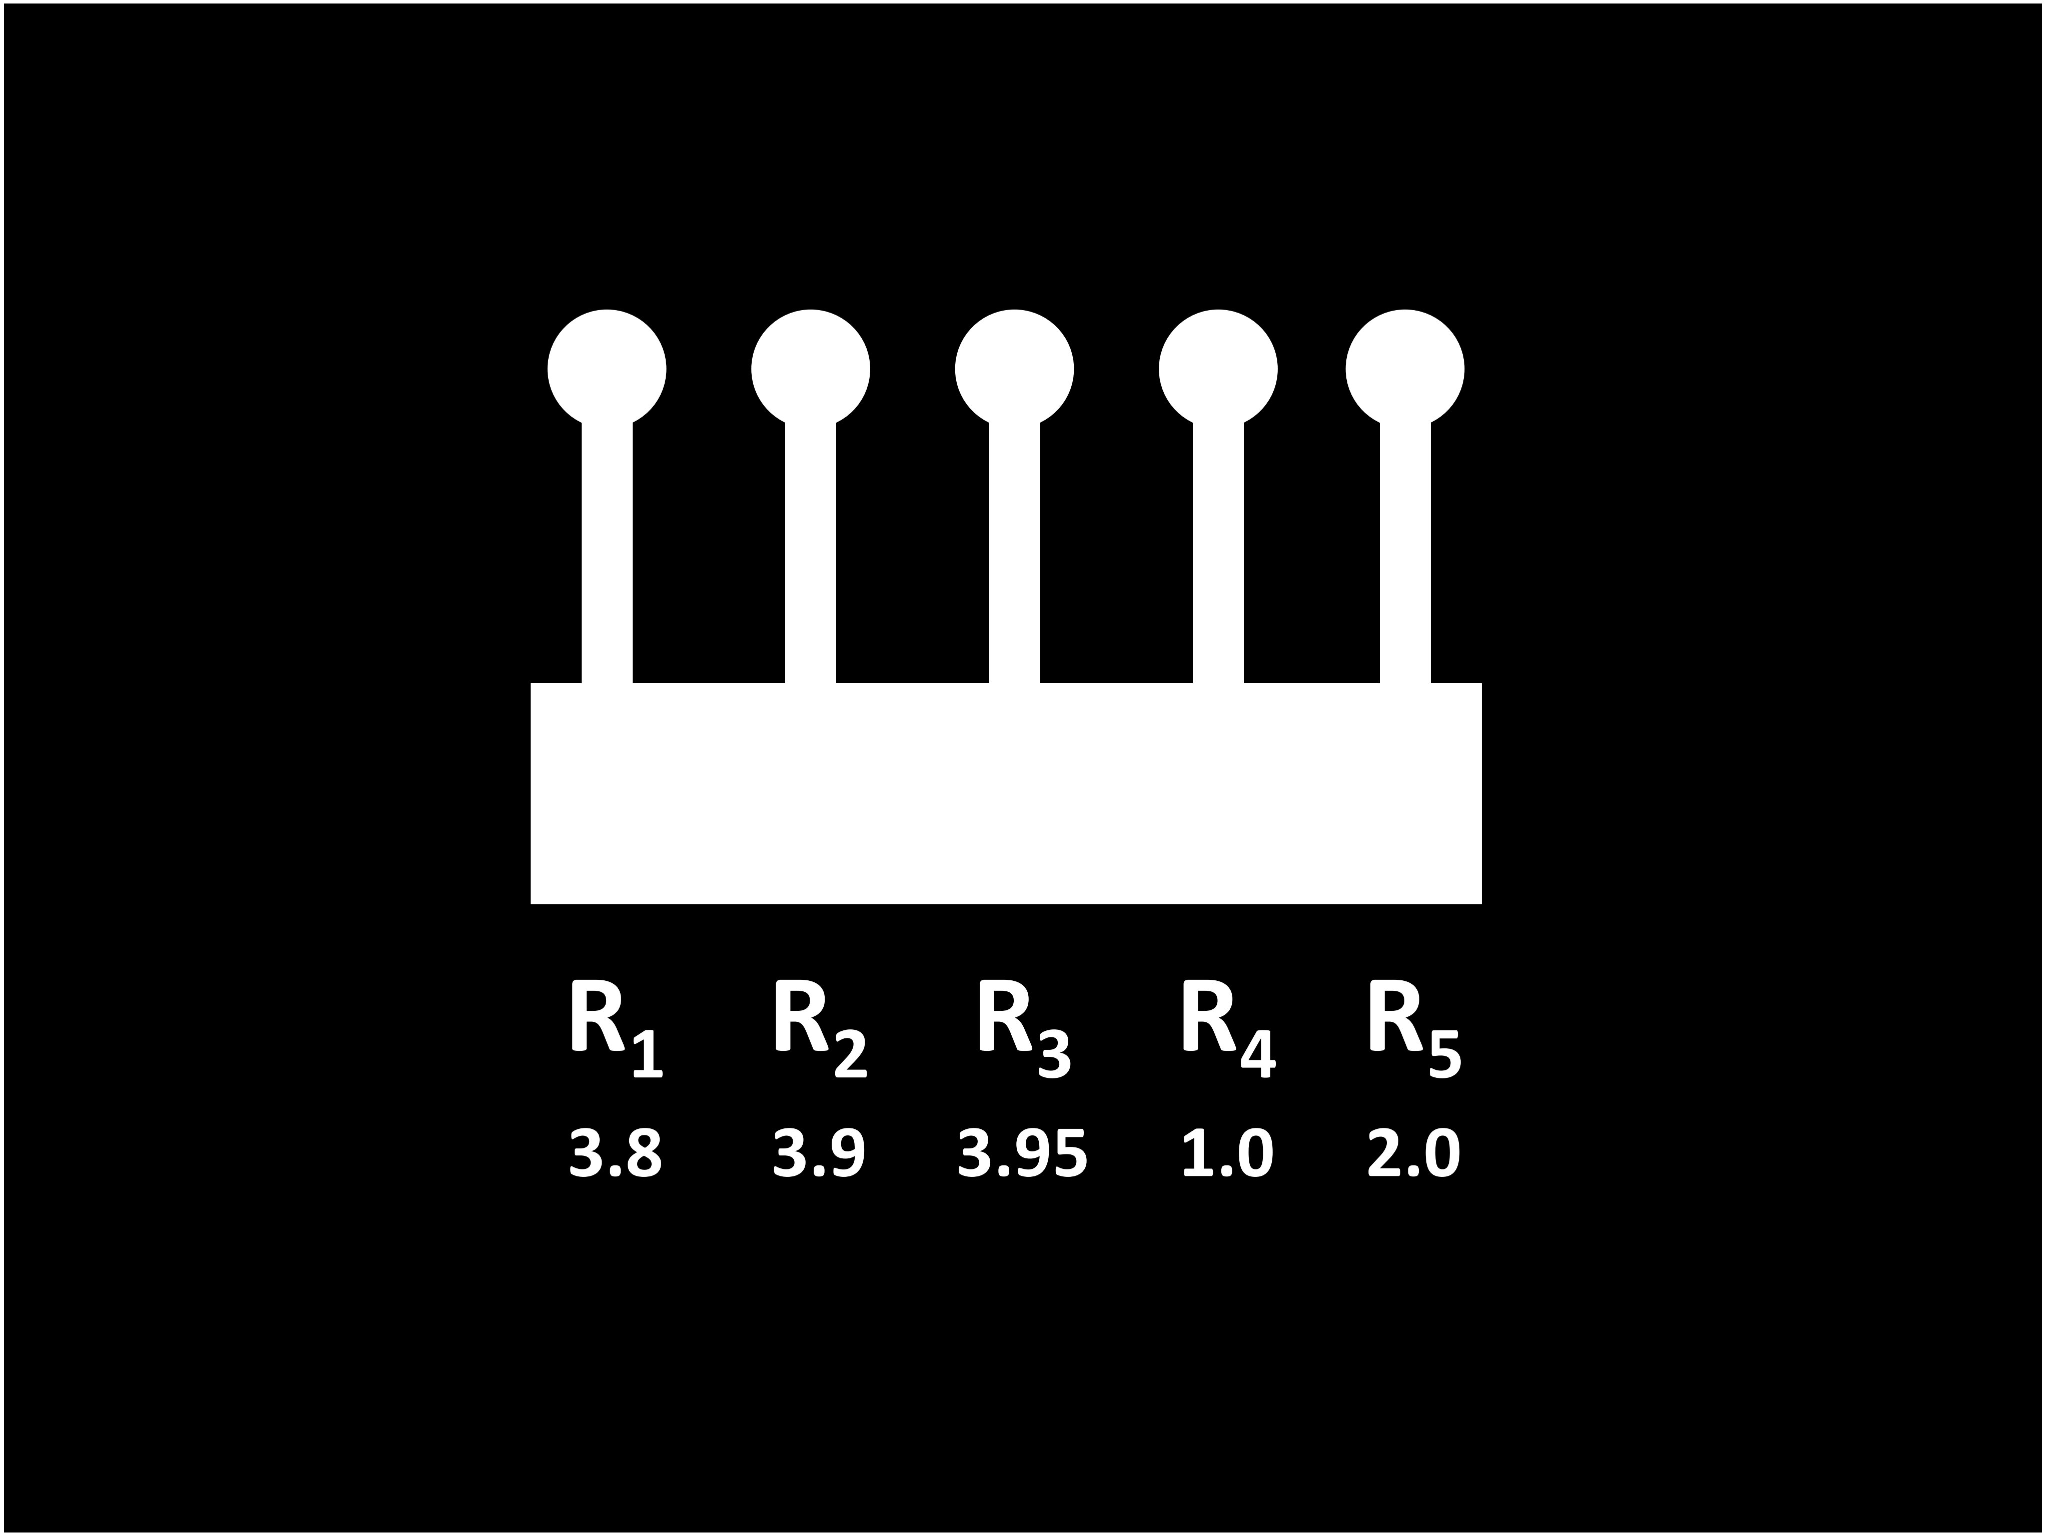
R_1 R_2 R_3 R_4 R_5
3.8 3.9 3.95 1.0 2.0

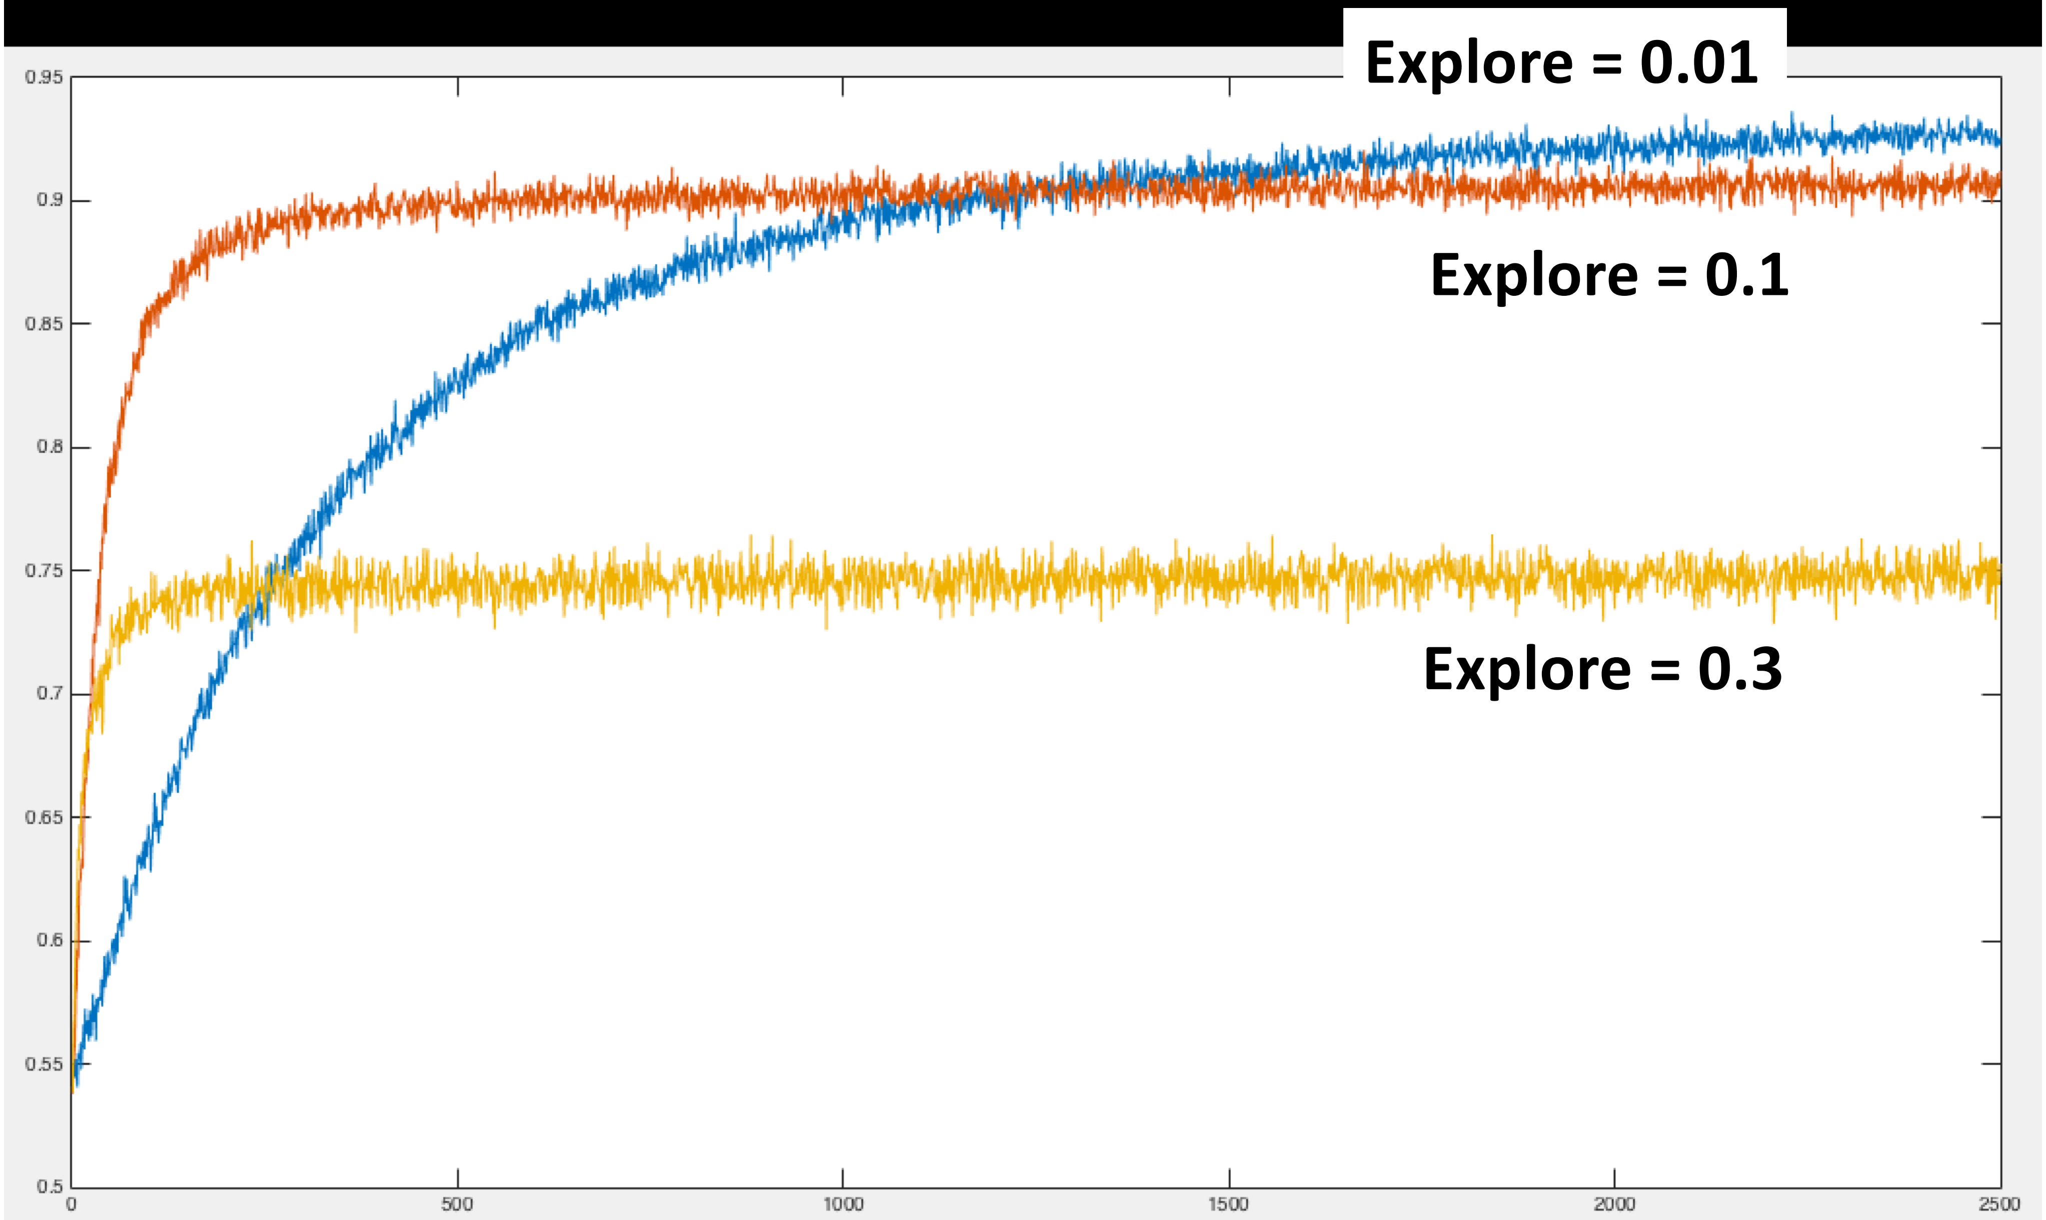
Explore = 0.01
Explore = 0.1
Explore = 0.3
0.95
0.9
0.85
0.8
0.75
0.7
0.65
0.6
0.55
0.5
0
500
1000
1500
2000
2500

Why Explore?

1. Unknown Values

2. Changing Environments

Should exploration rates change over time?

Decision Making

1. Always Choose the Highest Value Option
2. Exploration versus Exploitation

BUT

3. Psychological factors have to be accounted for