CS 112: Modeling Uncertainty in Information Systems

Prof. Jenn Wortman Vaughan

June 4, 2012

Lecture 17

Reminders & Announcements

- Homework 5 is due in section this Friday, June 8
- Course evaluations are available for you to complete online – we'll end a little early on Thursday in case you want to use that time to complete them
- Friday's section will be a review session for the final, so bring any questions that you have

Final Exam

- Tuesday, June 12, 11:30am – 2:30pm in this room
- One double-sided sheet of hand-written notes allowed
- No other notes, books, calculators, cell phones, etc.
- Best way to study is to practice solving problems like the end-of-chapter problems in the book (solutions available on the book's website)

You should be able to…

- Figure out when and how to apply the basic rules and definitions of probability, including
  - the multiplication rule
  - Bayes' rule
  - the total probability theorem
  - the total expectation theorem
  - the definition of conditional probability
  - the definition of independence
  - the counting principle
  - the definitions of expectation and variance

You should be able to…

- Translate word problems into math
  - Define the appropriate random variables or events
  - Be able to state the quantity that you must solve for in terms of these random variables or events
  - Recognize common random variables (Bernoulli, binomial, and geometric) in problem descriptions

Topics Covered

- Sample spaces, events, and basic probability (Chapter 1)
- Discrete and continuous random variables (Chapters 2–3)
- Covariance and correlation (4.2)
- Markov, Chebyshev, law of large numbers (5.1, 5.2)
- Information, entropy, coding (as covered in class)
- MAP and maximum likelihood (as covered in class)
- Naive Bayes (external reading linked from website)
- Markov chains (7.1–7.4; see assignments on website)

You are responsible for any material that was covered in the reading, in class (including this week), or in section!

Back to Markov Chains…

Steady-State Convergence Theorem

Theorem: Consider any Markov chain with a single recurrent class, which is not periodic. There are unique values $\pi_1, \ldots, \pi_m$ that satisfy the balance equations:

$$\pi_j = \sum_{k=1}^{m} \pi_k p_{k,j} \qquad \text{for } j = 1, \ldots, m$$

$$\sum_{k=1}^{m} \pi_k = 1$$

and for each j, π_j is the long term fraction of time that the state is j. These are called the steady-state probabilities.

Convergence of Markov Chains

- What if there is only one recurrent class but it is periodic?

Convergence of Markov Chains

- What if there is only one recurrent class but it is periodic?
  - Limit of $P(X_n = i)$ is not well defined, but balance equations still give us the long term frequencies

Convergence of Markov Chains

- What if there is only one recurrent class but it is periodic?
  - Limit of $P(X_n = i)$ is not well defined, but balance equations still give us the long term frequencies
- What if there are multiple recurrent states?

Convergence of Markov Chains

- What if there is only one recurrent class but it is periodic?
  - Limit of $P(X_n = i)$ is not well defined, but balance equations still give us the long term frequencies
- What if there are multiple recurrent states?
  - Once a recurrent state is entered, the same ideas apply

Convergence of Markov Chains

- What if there is only one recurrent class but it is periodic?
  - Limit of $P(X_n = i)$ is not well defined, but balance equations still give us the long term frequencies
- What if there are multiple recurrent states?
  - Once a recurrent state is entered, the same ideas apply
  - How likely are we to enter each recurrent state?

Absorption

The Gambler's Ruin

A gambler plays rounds of games in a casino. At each round, he wins \$1 with probability p and loses \$1 with probability 1-p, independent of what happens in other rounds. He continues playing until he either accumulates a target amount of money \$$m$, or loses all of his money.

- How can we represent this scenario as a Markov chain?

The Gambler's Ruin

A gambler plays rounds of games in a casino. At each round, he wins \$1 with probability p and loses \$1 with probability 1-p, independent of what happens in other rounds. He continues playing until he either accumulates a target amount of money \$$m$, or loses all of his money.

- How can we represent this scenario as a Markov chain?
- What are the recurrent classes?

The Gambler's Ruin

A gambler plays rounds of games in a casino. At each round, he wins \$1 with probability p and loses \$1 with probability 1-p, independent of what happens in other rounds. He continues playing until he either accumulates a target amount of money \$$m$, or loses all of his money.

- How can we represent this scenario as a Markov chain?
- What are the recurrent classes?
- What is the probability that the gambler accumulates \$$m$?

Absorbing States

- An absorbing state in a Markov chain is a single-state recurrent class — that is, any state i such that $p_{i,i}=1$

Absorbing States

- An absorbing state in a Markov chain is a single-state recurrent class — that is, any state i such that $p_{i,i}=1$

- Assume we have a Markov chain in which every state is either absorbing or transient

Absorbing States

- An absorbing state in a Markov chain is a single-state recurrent class — that is, any state i such that $p_{i,i}=1$

- Assume we have a Markov chain in which every state is either absorbing or transient
  - Let s be a particular absorbing state

Absorbing States

- An absorbing state in a Markov chain is a single-state recurrent class — that is, any state i such that $p_{i,i}=1$

- Assume we have a Markov chain in which every state is either absorbing or transient
  - Let s be a particular absorbing state
  - Let E be the event that $X_n = s$ for *any* time step $n \geq 0$

Absorbing States

- An absorbing state in a Markov chain is a single-state recurrent class — that is, any state i such that $p_{i,i} = 1$

- Assume we have a Markov chain in which every state is either absorbing or transient
  - Let s be a particular absorbing state
  - Let E be the event that $X_n = s$ for *any* time step $n \geq 0$
  - For each state i, let $a_i = P(E \mid X_0 = i)$

Absorbing States

- An absorbing state in a Markov chain is a single-state recurrent class — that is, any state i such that $p_{i,i}=1$

- Assume we have a Markov chain in which every state is either absorbing or transient
  - Let s be a particular absorbing state
  - Let E be the event that $X_n = s$ for *any* time step $n \geq 0$
  - For each state i, let $a_i = P(E \mid X_0 = i)$

What do we know about the values a_i?

Absorption Probability Equations

Theorem: Consider a Markov chain in which each state is either absorbing or transient. Fix a particular absorbing state s. Then the probabilities a_i of eventually reaching state s after starting at state i are the unique solutions to the following system of equations:

$$a_s = 1$$

$$a_i = 0 \qquad \text{for all absorbing } i \neq s$$

$$a_i = \sum_{j=1}^{m} p_{i,j} a_j \qquad \text{for all transient } i$$

The Gambler's Ruin

A gambler plays rounds of games in a casino. At each round, he wins \$1 with probability p and loses \$1 with probability 1-p, independent of what happens in other rounds. He continues playing until he either accumulates a target amount of money \$$m$, or loses all of his money.

- What is the probability that the gambler accumulates \$$m$?

The Gambler's Ruin

A gambler plays rounds of games in a casino. At each round, he wins \$1 with probability p and loses \$1 with probability 1-p, independent of what happens in other rounds. He continues playing until he either accumulates a target amount of money \$$m$, or loses all of his money.

- What is the probability that the gambler accumulates \$$m$?

(we'll just do the case where $p = .5$, but you can read the general solution in example 7.11 in the text)

Gambler's Ruin

$m = 1000 \qquad p = 0.4999$

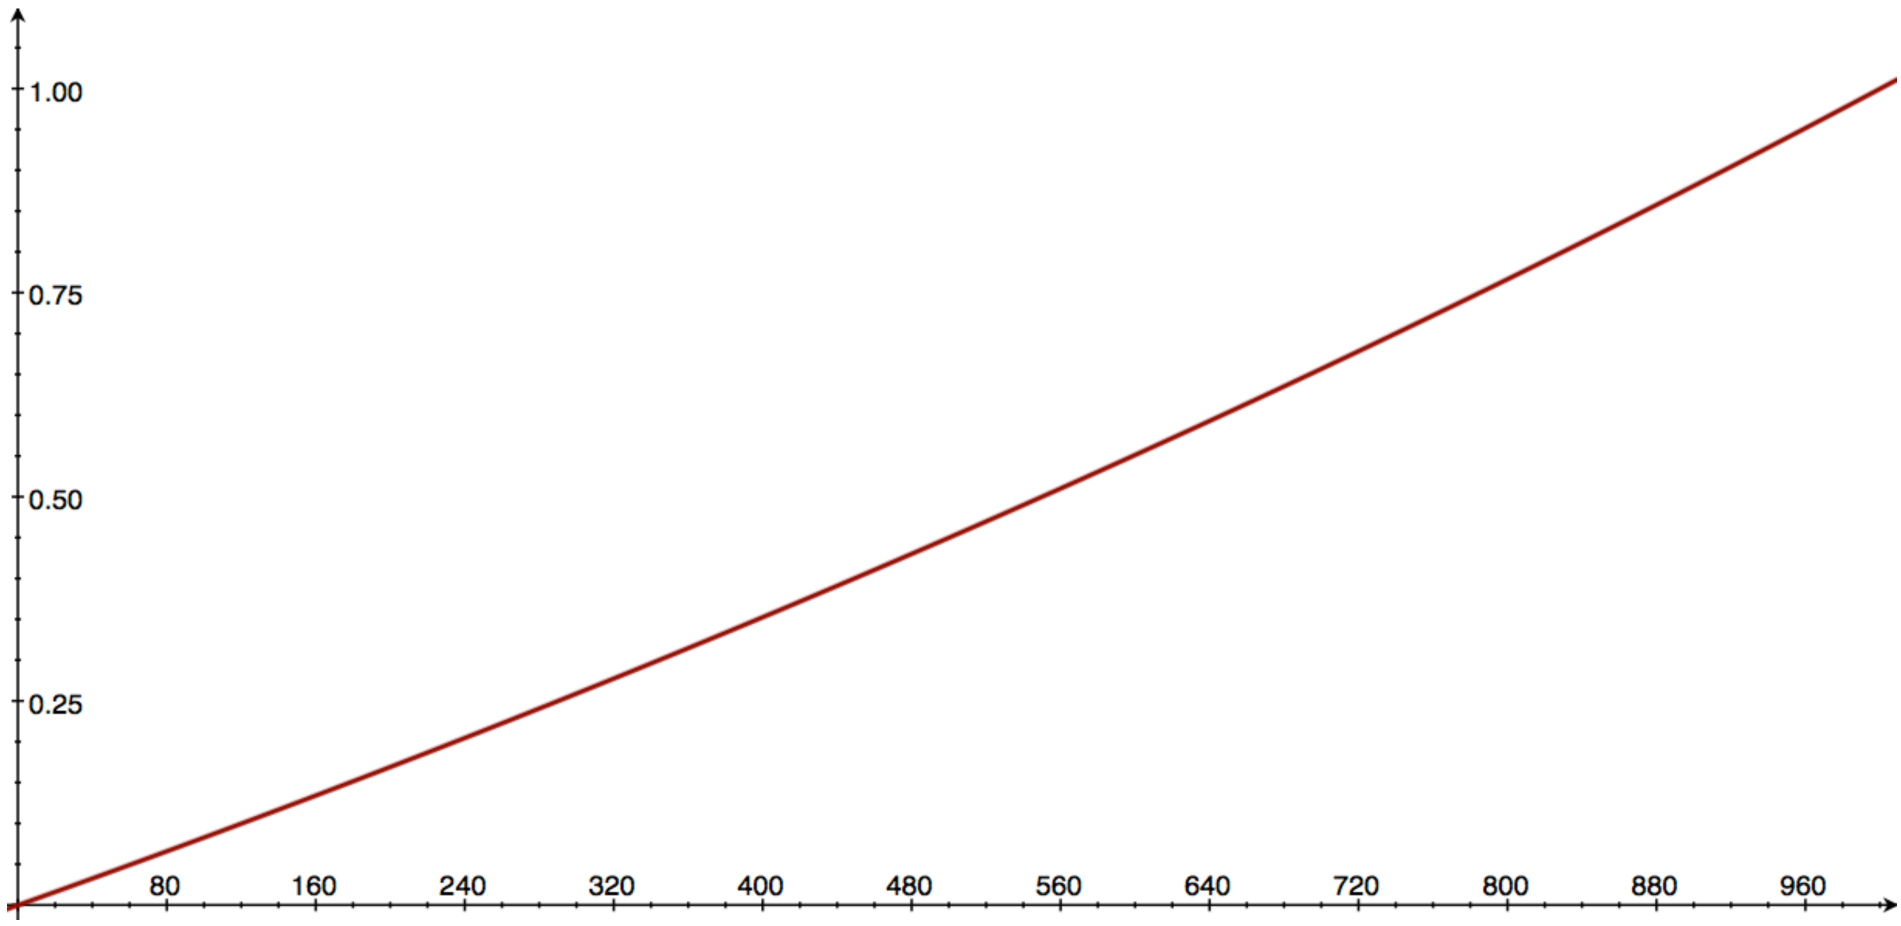

Gambler's Ruin

$m = 1000 \qquad p = 0.499$

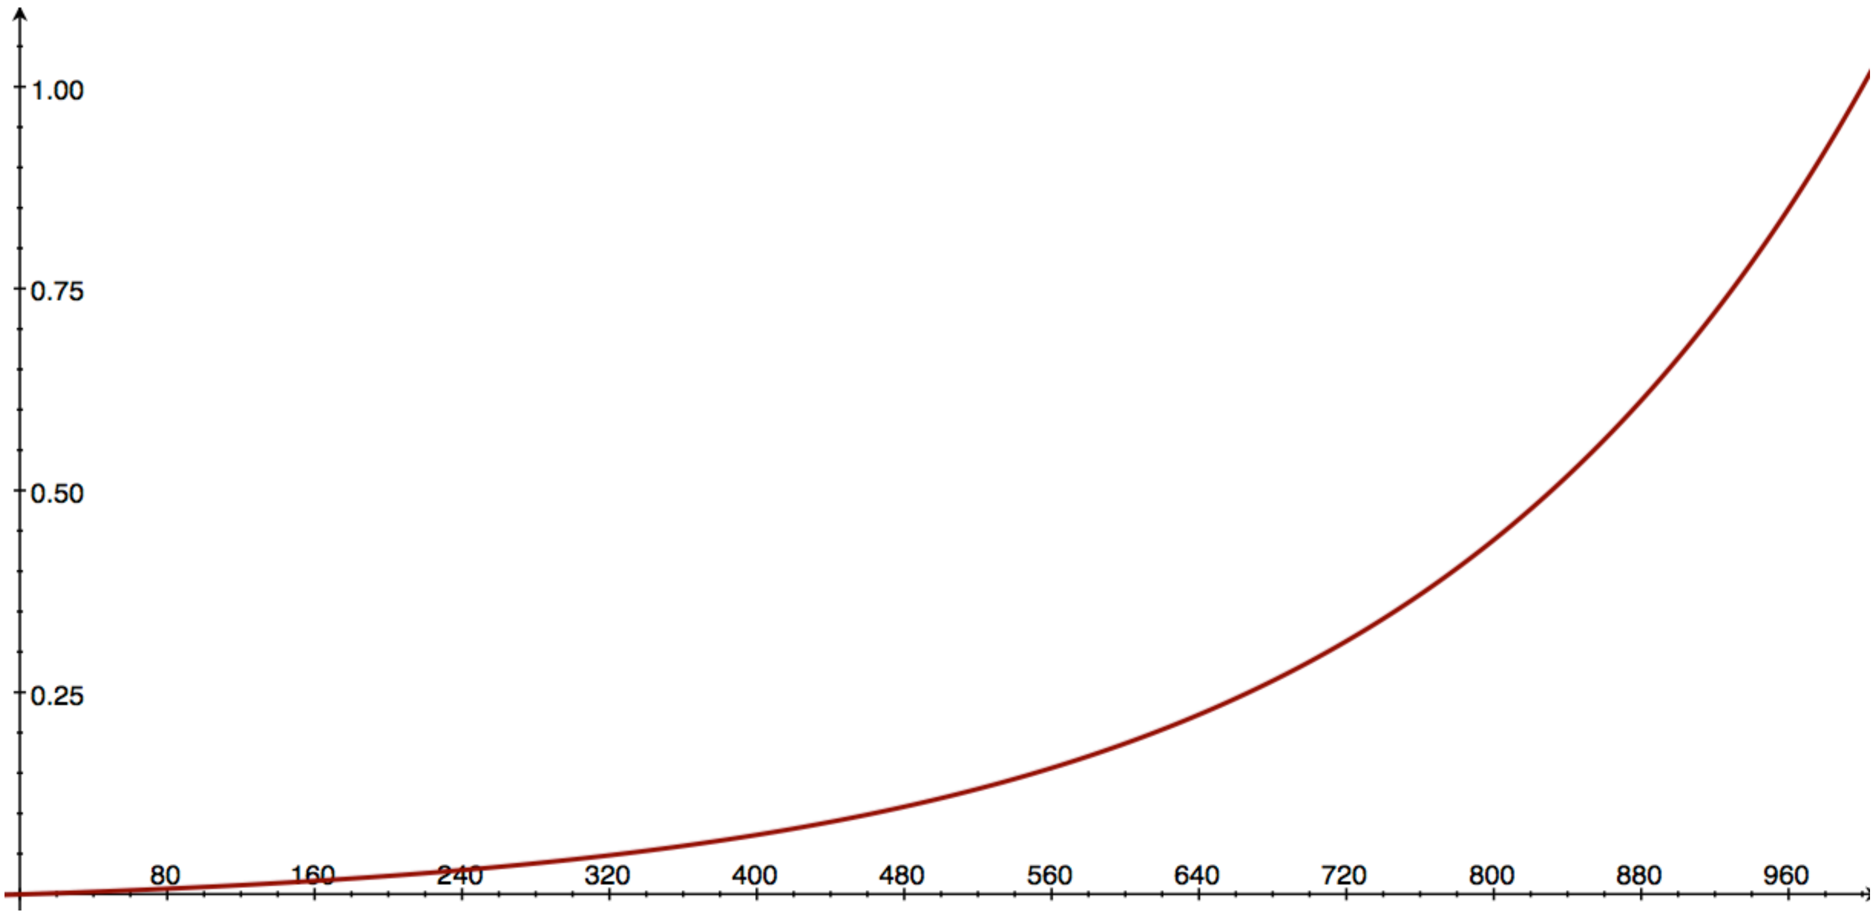

Gambler's Ruin

$m = 1000 \quad p = 0.49$

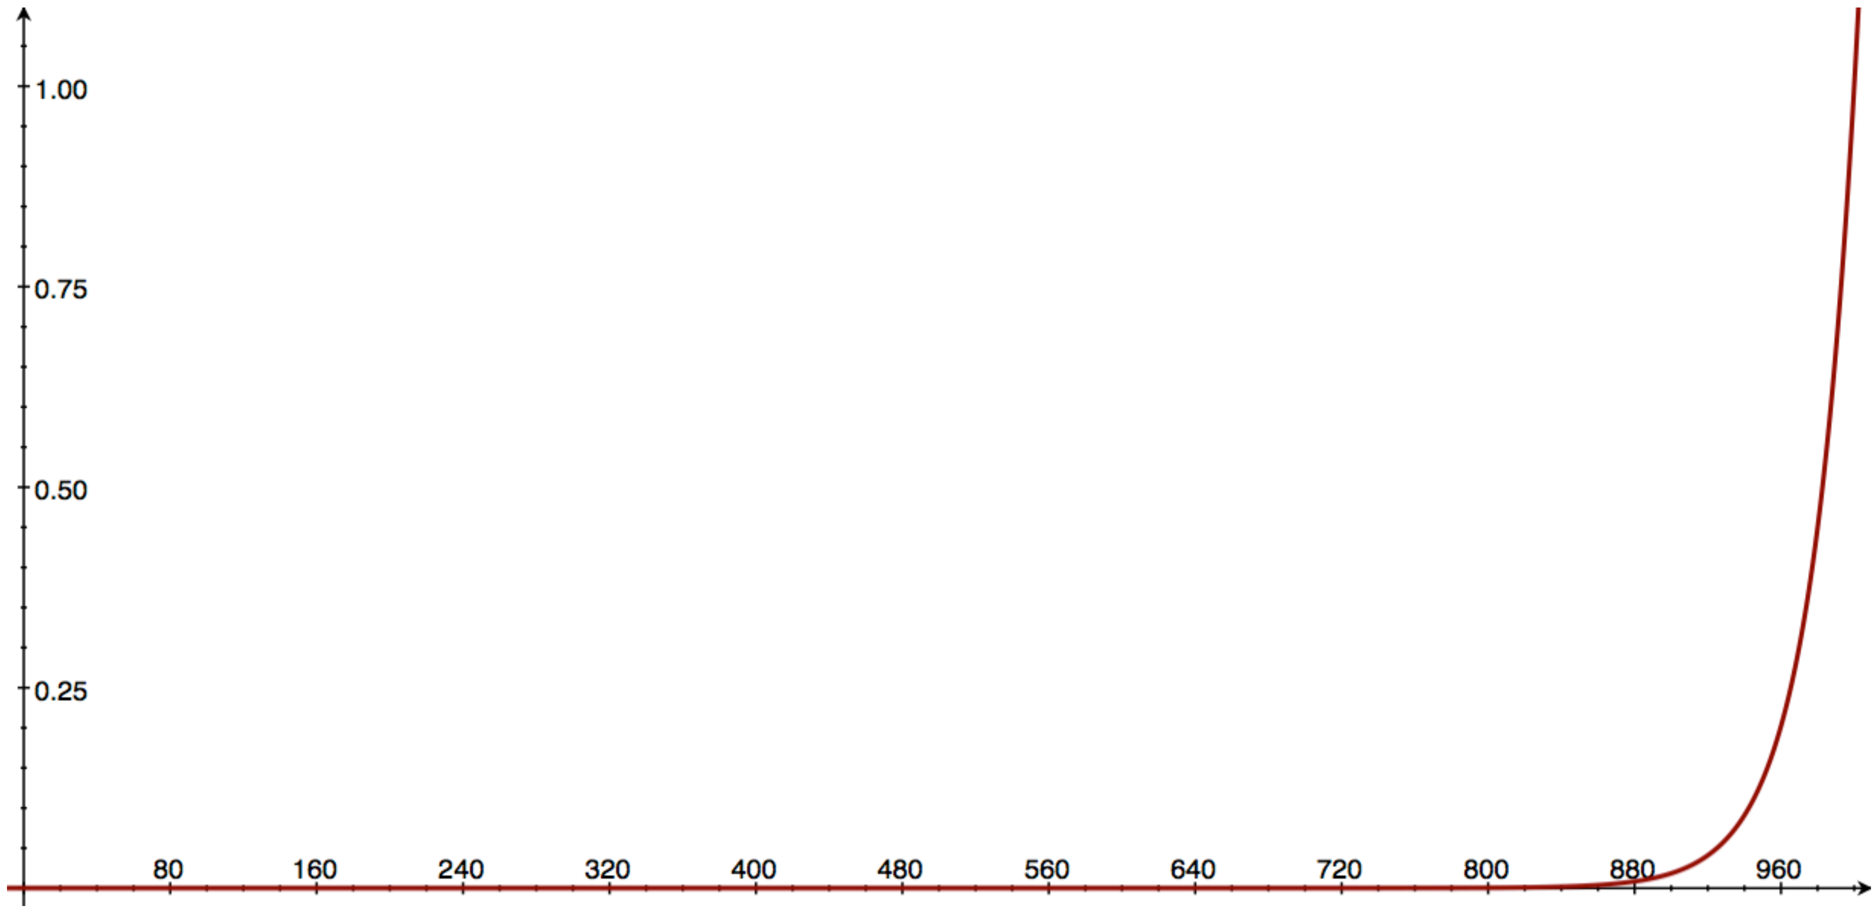

Gambler's Ruin

$m = 1000 \quad p = 0.4$