COS 445 — Final

Due online Wednesday, May 15th at 11:59 pm

- All problems on this exam are *no collaboration* problems.
- You may not discuss any aspect of any problems with anyone except for the course staff.
- You may not consult any external resources, the Internet, etc.
- You *may* consult the course lecture notes on Ed, any of the five course readings, past Ed discussion, or any notes directly linked on the course webpage (e.g. the cheatsheet, or notes on linear programming).
- You *may* discuss the test with the course staff, but we will only answer clarification questions and will not give any guidance or hints. You should feel free to ask any questions and let us judge whether or not to answer, but just know that we may choose to politely decline to answer. We may choose to answer questions with a response of "I'm sorry, but I'm not comfortable answering that question," or "it is within the scope of the exam for you to answer that question yourself" (or some variant of these).
- If you choose to ask a question on Ed, *ask it privately*. We will maintain a pinned FAQ for questions that are asked multiple times (please also reference this FAQ).
- Please upload each problem as a separate file via codePost, as usual.
- You *may not* use late days on the exam. You must upload your solution by May 15th at 11:59pm. If you are working down to the wire, upload your partial progress in advance. **There is no grace period for the exam**. In case of a true emergency where you cannot upload, email both smweinberg@princeton.edu and dallagnol@princeton.edu your solutions asap.
- If you miss the deadline, **even by a minute**, university policy prohibits us from grading your exam without explicit permission from your dean. Please make sure you have something submitted by the deadline, and take into account that the server may be overloaded or sluggish near the end. For example, you may wish to treat 11:45pm on May 15th as a "pencils down", to leave yourself enough time to safely upload (or email your solutions if codePost crashes).
- There are *no exceptions, extensions, etc.* to the exam policy (again, in case of a *truly exceptional* circumstance, you should reach out to your residential dean and have them contact us).

Problem 1: COS 445 Speedrun (120 points)

For each of the 12 problems below: **unless otherwise specified**¹ you do not need to show any work and can just state the answer. However, if you simply state an incorrect answer with no justification, we cannot award partial credit. You are encouraged to provide a **very brief** outline/justification in order to receive partial credit in the event of a tiny mistake. For example, we will award very significant partial credit if you clearly execute the correct outline, but make a mistake in implementation.

Part a: Stable Matchings (10 points)

Four students, Alice, Bob, Claire and David are applying to summer internships at Apple, Bell Labs, Capital One and Dell (all of which need exactly one intern). Here are their preferences, sorted from favorite to least favorite:

- Alice: Apple \succ Dell \succ Bell Labs \succ Capital One.
- Bob: Apple \succ Bell Labs \succ Dell \succ Capital One.
- Claire: Capital One \succ Dell \succ Bell Labs \succ Apple.
- David: Bell Labs \succ Dell \succ Apple \succ Capital One.

and the companies preferences:

- Apple: Alice \succ Claire \succ Bob \succ David.
- Bell Labs: Bob \succ Alice \succ Claire \succ David.
- Capital One: Alice \succ David \succ Claire \succ Bob.
- Dell: Alice \succ Bob \succ Claire \succ David.

Find the stable matching that arises from Deferred Acceptance when the students propose. A reminder of the Deferred Acceptance algorithm is the Lecture Stable Matchings I.

Part b: Voting Rules (10 points)

A town of 20 voters is holding an election between candidates Alice, Bob and Carol.

- 9 of the voters prefer Alice \succ Bob \succ Carol
- 6 of the voters prefer Carol \succ Bob \succ Alice
- 5 of the voters prefer $Bob \succ Carol \succ Alice$

State the winning candidate selected by each of the following voting rules: Borda, IRV, Plurality. A reminder of these three voting rules is in Lecture Voting Theory I.

¹If otherwise specified, you should follow the otherwise specifications.

Part c: Game Theory (10 points)

Find a Nash equilibrium of the following game and state the expected payoff for both players. A definition of Nash equilibrium can be found in Lecture Game Theory II.

Player X, the row player, chooses between actions x_1 and x_2 . Player Y, the column player, chooses between actions y_1 and y_2 . The first number in each box denotes the payoff to X, and the second number is the payoff to Y. For example, if X plays action x_1 and the column player plays action y_1 , then X gets payoff 5 and Y gets payoff 5.

	y_1	y_2
x_1	(5,5)	(-1,8)
x_2	(0,5)	(4,4)

Part d: Extensive Form Games (10 points)

Consider the extensive form game in Figure 1.



Figure 1: An extensive form game.

There are two players (named 1 and 2) and three rounds. First player 1 plays, then player 2, then player 1 again. The numbers on the leaves denote the payoffs to the first and second players, respectively (as labeled at the internal nodes of the tree). The labels on the edges denote the names of the actions they can play at that turn.

(i) Find a subgame-perfect Nash equilibrium for this game.

(ii) Find a pure Nash equilibrium such that both players receive strictly higher payoff than in the subgame-perfect Nash equilibrium from (i).

Recall that for both parts, a complete strategy lists a pure action for **every** internal node. For example, $\{L, S\}$ is not a complete strategy for player 1, nor is *a* a complete strategy for player 2. But $\{L, S, U, W, Y\}$ is a complete strategy for player 1, as is $\{a, c\}$ for player 2. A definition of Nash equilibria and subgame perfect Nash equilbria can be found in Lecture Game Theory II.

Part e: Linear Programming (10 points)

Write the dual of the following LP. You do not need to solve the LP. You only need to write the dual. A reminder of LP duality is in Lecture Linear Programming.

Maximize 2x + 5y, such that:

- $1x + 6y \le 2$.
- $7x + 3y \le 2$.
- $x, y \ge 0$.

Part f: Scoring Rules (10 points)

Suppose you are asked to predict tomorrow's weather. There's four possible outcomes: it will be sunny, rainy, cloudy or snowy. You will be paid according to the logarithmic scoring rule $(S(\vec{x}, i) = \log_2(x_i))$. What is your expected payoff if you report the uniform distribution over these four outcomes?²

A reminder of notation for scoring rules is in Lecture Scoring Rules.

Part g: Welfare-maximizing Auctions (10 points)

There are three bidders, and two ad slots. The first ad slot has a click-through rate of 1, and the second has a click-through rate of 1/3. The three bidders submit bids of $b_1 = 7, b_2 = 4, b_3 = 2$. The auctioneer is running a VCG auction (to assign each bidder at most one slot, and each slot to at most one bidder).

For each of the three bidders, state the slot they win and their payment (state the bidder's total payment, not their payment per-click).

A reminder of the VCG auction for sponsored search is in Lecture Auction Theory II.

Part h: Revenue-maximizing Auctions

Suppose you are selling a pen to a single buyer. The buyer's value is drawn uniformly from [4, 20]. What is the revenue-optimal auction (menu) for you to sell the pen? Also, state the expected revenue you achieve.

Note: You may use without proof that the PDF of the uniform distribution on [4, 20] is equal to $\frac{1}{20-4}$ on the entire interval [4, 20]. You may also use without proof that the CDF $F(\cdot)$ of the uniform distribution on [4, 20] satisfies $F(x) = \frac{x-4}{20-4}$ when $x \in [4, 20]$.

Single-bidder revenue-maximizing auctions are computed for examples in Lecture Auction Theory III.

²Note that the problem does not tell you your true belief, and this information is not necessary to solve the problem.

Part i: Price of Anarchy (10 points)

Consider the network in Figure 2. There are two nodes, s and t, and one unit of flow traveling from s to t. There are two directed edges from s to t, one with cost c(x) = 1 and the other with cost $c(x) = x^2$. Compute the Price of Anarchy of this instance.

A reminder of Price of Anarchy appears in Lecture Price of Anarchy I.



Figure 2: A routing network.

Part j: Cake cutting (10 points)

There is a single cake, the unit-interval [0, 1]. Alice, Bob, and Charlie all have normalized, additive valuations (that is, v([0, 1]) = 1, $v(\emptyset) = 0$, and $v(X \cup Y) = v(X) + v(Y)$ whenever $X \cap Y = \emptyset$, and $v(X) \ge 0$ for all X). Alice's valuation satisfies $v_A([0, 1/5]) = 1$, distributed uniformly. Bob's satisfies $v_B([1/2, 5/6]) = 1$, distributed uniformly. Charlie's satisfies $v_C([1/2, 1]) = 1$, distributed uniformly.

Consider the allocation which awards Alice the interval [0, 1/3], Bob the interval [1/3, 2/3], and Charlie the interval [2/3, 1]. Is the allocation proportional? Is it envy-free? Is it equitable?

A reminder of these terms appears in Lecture Fair Division and Cake Cutting.

Part k: Behavioral Economics (10 points)

Recall that a utility function $f(\cdot)$ takes as input a deterministic outcome and outputs a utility in \mathbb{R} . Say that there are three possible deterministic outcomes, A, B, C. **Define a utility function** $f(\cdot)$ **such that** an expected utility maximizer with utility function $f(\cdot)$ prefers the randomized outcome which is A with probability 1/3, B with probability 1/3, and C with probability 1/3 to the randomized outcome which is A with probability 1/5, B with probability 2/5, and C with probability 2/5.

A reminder of expected utility maximizers appears in Lecture Behavioral Game Theory I.

Part *l***:** Time-Inconsistent Planning (10 points)

In the planning graph of Figure 3:

- What is the shortest path from s to t?
- What path is taken by a naive planner with present bias b = 2?

• What path is taken by a sophisticated planner with present bias b = 2?

A reminder the naive planner and sophisticated planner is in Lecture Behavioral Game Theory II.



Figure 3: A planning graph. Edges are color-coded to reduce confusion.

Problem 2: Advertising for Auctions (55 points)

Recall the following definition, from the Cheatsheet Section 1.8. You may wish to refer to Cheatsheet Section 1.8 for further help parsing the definition. It is possible to fully solve this problem without visiting the Cheatsheet. However, once you figure out your solution, the Cheatsheet may help you write a simpler proof with significantly fewer calculations.

Definition 1 (Stochastic Dominance). We say that a single-variable distribution D^+ stochastically dominates D if for all x, $\Pr_{v^+ \leftarrow D^+}[v^+ \ge x] \ge \Pr_{v \leftarrow D}[v \ge x]$. Put another way, if F^+ is the CDF of D^+ , and F is the CDF of D, then $F^+(x) \le F(x)$ for all x.

Below, let $\operatorname{REV}(D) := \max_{x} \{x \cdot \operatorname{Pr}_{v \leftarrow D}[v \ge x]\}$ denote the expected revenue of the optimal auction for selling to a single buyer whose valuation is drawn from D. You may assume that $\operatorname{REV}(D)$ is well-defined. Let $p(D) := \arg \max_{x} \{x \cdot \operatorname{Pr}_{v \leftarrow D}[v \ge x]\}$ denote the optimal price to set (break ties in favor of the minimum price in the $\arg \max$).³ Finally, let $\operatorname{REV}_p(D_1, \ldots, D_n)$ denote the expected revenue of the second-price auction with reserve p when there are n bidders with valuations drawn independently from D_1, \ldots, D_n .

Part a (15 points)

Let D^+ stochastically dominate D. Prove the following statement or find a counterexample: $\operatorname{REV}(D^+) \geq \operatorname{REV}(D)$.

For Part a, if you choose to find a counterexample, it can be continuous or discrete.

Part b (15 points)

Let D^+ stochastically dominate D. Prove the following statement or find a counterexample: $p(D^+) \ge p(D)$.

For Part b, if you choose to find a counterexample, it can be continuous or discrete.

Part c (25 points)

Prove the following statement or find a counterexample: For all reserve prices $p \ge 0$, and all number of bidders n, if D_i^+ stochastically dominates D_i for all $i \in [n]$, then $\operatorname{Rev}_p(D_1^+, \ldots, D_n^+) \ge \operatorname{Rev}_p(D_1, \ldots, D_n)$.⁴

For Part c, if you choose to find a counterexample, it can be continuous or discrete.

³Note that this is always well-defined, as long as Rev(D) is well-defined.

⁴If you choose to prove this statement, it must hold for all n and all p. If you choose to find a counterexample, observe that you just need to find a single $p \ge 0$ and $n \ge 1$ and instance $D_1, \ldots, D_n, D_1^+, \ldots, D_n^+$.

Problem 3: Silly Selfish Mining (45 points)

If you need a refresher, refer to Lecture Bitcoin Mining Game and Selfish Mining that this problem is based on. That game is motivated in Lecture Bitcoin Intro (but it is most clearly stated in Lecture Bitcoin Mining Game and Selfish Mining, so we suggest going straight there).

Imagine that you control an α fraction of the total computational power in the Bitcoin network, all other miners use longest-chain and **always tie-break against you**,⁵ and you use the following mining strategy.⁶ To be explicit, you will use the strategy Silly Selfish Mining, and the rest of the network uses the strategy Honest (both are defined below).

We will use the following notation/definitions, which are copied from Lecture Bitcoin Mining Game and Selfish Mining, Section 4:

- The **height** of a block is equal to the length of the path from it to the root. For example, a block that points directly to the root has height 1. If B points to a block of height h, then B has height h + 1.
- G refers to the set of blocks that have been announced so far.
- G'_m refers to the set of blocks that were created by m.
- $G_m := G \cup G'_m$ refers to the set of blocks that are either announced so far, or created by m.

Recall that the order of operations and notation during a round t are (you are m):

- a) A miner is selected to create a block, equal to you with probability α , and not equal to you with probability 1α .
- b) All miners $\neq m$ announce any blocks they want, which immediately adds these blocks to G.
 - b.i) For use of notation below, we let h(t) denote the maximum height among all blocks in G, at this point.
 - b.ii) For use of notation below, we let $h_m(t)$ denote the maximum height among all blocks in G'_m (that is, the maximum height among all blocks that are created by m).
- c) Miner m announces any blocks they want, and this immediately adds these blocks to G.

Below are the two relevant strategies for this problem:

Silly Selfish Mining:

- At every time step t, mine on top of the longest chain in G_m (if you are selected to mine), tie-breaking in favor of your own blocks. That is, you will mine on top of the longest chain among blocks that were announced by others, or that you created yourself (and may or may not have announced).
- **Do not announce** your blocks immediately upon mining them.

⁵Note that this is the same setup as in Section 4 of Lecture Bitcoin Mining Game and Selfish Mining. In Lecture Bitcoin Mining Game and Selfish Mining Section 3, we made a temporary strong assumption that miners tiebreak in favor of you.

⁶Note that it is the same as Selfish Mining from Lecture Bitcoin Mining Game and Selfish Mining Section 4, *except there is no special case*. The entire setup is the same as in Lecture Bitcoin Mining Game and Selfish Mining, except the strategy below does not have the special case of Selfish Mining.

- During timestep *t*, choose blocks to announce according to the following rules:
 - i) If another miner announced a block of height h(t) during round t, and $h_m(t) = h(t)$, announce your block of height h(t).
 - ii) If another miner announced a block of height h(t) during round t, and $h_m(t) > h(t)+1$, announce your block of height h(t).
 - iii) If another miner announced a block of height h(t) during round t, and $h_m(t) = h(t)+1$, announce your two blocks of height h(t) and h(t) + 1.

Honest:

- At every time step t, mine on top of the longest chain in G (if you are selected to mine), tie-breaking in favor blocks not created by m. That is, Honest miners will mine on top of the longest chain among blocks that were previously announced, and will tie-break against m.
- Announce your blocks immediately upon mining them.

Part a (10 points)

Consider the case when miners are selected according to the following sequence $(\neq m \text{ is an honest miner})$: $\langle M_1, M_2, M_3, M_4, M_5, M_6, M_7, M_8 \rangle = \langle m, \neq m, m, m, m, \neq m, \neq m, \neq m \rangle$. What is h(t) and $h_m(t)$ for each t? You should present your answer in the following format: $\langle h(1), h(2), \ldots, h(8) \rangle$, $\langle h_m(1), h_m(2), \ldots, h_m(8) \rangle$, followed by a brief justification that each answer is computed correctly.⁷

Part b (10 points)

Present a sequence of selected miners $\langle M_1, M_2, M_3 \rangle$ (you must have exactly three time steps) where the actions taken by the silly selfish mining strategy above and the Selfish Mining Strategy presented in Lecture Bitcoin Mining Game and Selfish Mining differ in at least one timestep.⁸ Briefly justify why the two strategies behave differently.

Part c (25 points)

Describe (or draw, if you prefer) a Block Counting Scheme (as a function of α) to analyze the expected reward achieved with the silly selfish mining strategy defined above. You should also provide a brief explanation of why your analysis is correct. More specifically, you should provide:⁹

- 1. A Markov chain. Recall that a Markov chain requires:
 - A set V of states (V can be infinite).
 - A set *E* of transitions. Each $e = (u_e, v_e) \in E$ goes from one state u_e to another state v_e , and has a transition probability, q_e . Let E_u denote the set of transitions such that $u_e = u$. Then the transition probabilities must satisfy the following condition for every $u \in V$: $\sum_{e \in E_u} q_e = 1$.¹⁰

⁷If you want to be kind to the graders, you can also type |boxed| around your answer.

⁸That is, one strategy will decide to broadcast a block in a timestep in which the other does not.

⁹The specifications below just repeat the definition of a Block Counting Scheme from Lecture Bitcoin Mining Game and Selfish Mining.

¹⁰You are allowed to have multiple transitions between the same two nodes, if you choose to.

- 2. For all transitions $e \in E$, a variable H_e . H_e denotes the number of blocks for the honest miners (miners $\neq m$) that we would like to count as being in the steady-state longest chain when we use transition e.
- 3. For all transitions $e \in E$, a variable S_e . S_e denotes the number of blocks for the strategic miner (miner m) that we would like to count as being in the steady-state longest chain when we use transition e.
- 4. Every block that is in the steady-state longest chain should be counted during *exactly one* transition, and no other blocks should be counted. That is, every block in the steady-state longest chain must be counted during some transition, and it must never be double-counted in multiple transitions. Moreover, no block that is *not* in the steady-state longest chain can be counted during any transition.

You do not need to find the stationary probabilities of your Markov chain, prove that your Markov Chain has a steady-state distribution, nor compute the expected reward achieved by this strategy. **You only need** to describe the Block Counting Scheme as detailed above, and briefly explain why your Block Counting Scheme is correct.

Problem 4: Fair phone division (80 points)

You and your n - 1 best friends are deciding whether to buy a bundle deal on n new phones. In each of the following problems, there is a single bundle of n phones, and the total cost is B. Friend i has value $v_{ij} \ge 0$ for phone j.

- For all friends *i*, their valuation satisfies $\sum_{j} v_{ij} \ge B$ (their total value for all phones is at least *B*).
- Friend *i* gets payoff $v_{ij} p_i$ if they receive phone *j* and pay p_i . (This is true whether p_i is positive or negative). Each friend wants exactly one phone.
- Your job is to convince your friends to purchase the bundle **by designing an envy-free allo**cation of phones as well as how much each friend will pay. That is:
 - You must, taking as input v_{ij} for all i, j, propose which friend will receive which phone. Your procedure need not be strategyproof.
 - You must decide how much each friend will pay, and this sum must *exactly* cover the cost of the bundle: $\sum_i p_i = B$. Note that you are allowed to have $p_i < 0$ for some *i* if you want (think of this as one friend paying another to be OK with a junk phone). Your procedure need not be strategyproof.
 - Your final allocation/prices must be envy-free. That is, for all *i*, friend *i* must prefer their phone at the price they paid to any other phone *j* at the price paid by the friend who got phone *j*. To be extra formal, if *f*(*j*) denotes the friend who received phone *j*, and *h*(*i*) denotes the phone received by friend *i*, we must have that for all *i* and all *j*, *v*_{*ih*(*i*)} − *p*_{*i*} ≥ *v*_{*ij*} − *p*_{*f*(*j*)}.

Part a (10 points)

Prove that if an allocation/payment is envy-free (but might fail to satisfy $\sum_i p_i = B$), that for any $c \in \mathbb{R}$, updating the payments to $p'_i = p_i + c$ for all i (but keeping the allocation of phones exactly the same) is still envy-free.

Part b (10 points)

Prove that if an allocation/payment is envy-free, and $\sum_i p_i = B$, then $v_{ih(i)} - p_i \ge 0$ for all *i*. That is, prove that if the allocation/payment is envy-free, every player has non-negative payoff.

Part c (15 points)

Design a protocol to find an envy-free allocation/payment with $\sum_i p_i = B$ when n = 2, and prove that your protocol is correct.

Part d (20 points)

Design a protocol to find an envy-free allocation/payment with $\sum_i p_i = B$ when n = 3, and prove that your protocol is correct.

Hint: There is a solution which uses ideas from Lecture Fair Division and Cake Cutting (please note that this does *not* mean that the problem will be easy if you understand this lecture completely). There are other solutions as well.

Part e (25 points)

Design a protocol to find an envy-free allocation/payment with $\sum_i p_i = B$ when n = 4, and prove that your protocol is correct.

Hint: There is a solution which uses similar ideas to Part d, but there are other solutions as well.

Note: However you solve this problem, you may have to deal with case analysis. To make your solution readable for the grader, and for full credit, you should clearly state the important conclusions of your case analysis so that it is easy to follow.

Note: In case you do not find a full solution, you are encouraged (as usual) to write up *clearly-stated concrete partial progress*. Concrete partial progress for part e requires something beyond what is already solved in part d (but if you find an algorithm/analysis that correctly handles some special cases with 4 players, you are encouraged to write it up well for partial credit).

Note: The staff solution to part e is 1.5 pages, with spacious formatting and a verbose writing style. You are allowed to write as much as you like, but you may want to use this as a guide for the expected level of mathematical rigor. If you find yourself writing a ridiculous amount, and you see a clear mathematical subroutine that you would like to assume, it is OK to ask privately on Ed whether you may assume it (the answer may still be 'no').