Test 2 Solutions

Brendan Meeder

1 Nash

1.1

Playing the position (U,R) is a pure Nash equilibrium as neither player wants to change. To see whether there is a randomized equilibrium, we need to find the mixing rates p and q with which player A plays U and player B plays L, respectively. Player A is willing to randomize only if q is such that the expected payoff of playing U is the same as playing D. This gives us the equation

$$8q + 5(1-q) = 3q + 4(1-q) \iff 3q + 5 = 4 - q \iff q = -1/4.$$

At this point we know that there is no randomized strategy as q isn't in the interval $[0, 1]$. In order for a randomized strategy to exist, it must be the case that when we solve for p and q in this manner both values are in $[0, 1]$. It is *not* appropriate to 'clamp' the values to 0 or 1 if they are outside of these bounds.

1.2

Examining each of the four strategy tuples we see that there is no pure Nash equilibrium; in every position it is beneficial to at least one of the players to change their strategy. We now try to find a randomized equilibrium. As before, let p be the probability with which player A plays U and q be the probability with which B plays L. The condition under which A will randomize is

$$0 \cdot q + (-1)(1-q) = (-1)q + 2(1-q) \iff q - 1 = 2 - 3q \iff q = 3/4.$$

The condition under which B will randomize is

$$0 \cdot p + 1(1-p) = 1 \cdot p + (-2)(1-p) \iff 1 - p = 3p - 2 \iff p = 3/4.$$

Since both of these values are in $[0, 1]$ there is a randomized Nash equilibrium in which player A plays U with probability 3/4 and player B plays L with probability 3/4.

2 Second-price Auction Strategy

Part a

In a second-price auction for one item it is a dominant strategy to bid the truthful values. The winner is the bidder with the highest bid (in this case, with the highest valuation). So whoever bid $\max(A, B, C)$ wins and pays the second-largest value of $\{A, B, C\}$.

Part b

As we said in part a, bidders 1 and 2 will still bid their true valuations, regardless of how others are bidding.

Part c

In this case the expected payoff for bidder 1 will be less. First, consider the cases in which bidder 1 doesn't get item. If either bidder 2 or 3 bids a higher value than A, then the payoff for bidder 1 is zero. The probability that bidder 2's bid (the value B) is larger than A doesn't change based on bidder 3's new strategy, but the probability that bidder 3's bid of $(C+1)/2$ is bigger than A is now larger than the probability that $C > A$ as $(C+1)/2 \geq C$ on $[0,1]$, with equality occurring only at $C = 1$.

Now suppose that bidder 1 wins. Before there was an equal chance that 2's bid was bigger than 3's bid, and visa versa. Now, it's more likely that 3's bid of $(C+1)/2$ is bigger than B. In this case the payoff for bidder 1 is diminished, because it should would have been $\max(A-B, A-C) \geq \max(A-B, A-(C+1)/2)$. In the case that $B > (C+1)/2$, the payoff to 1 is unchanged.

We see that for a given outcome, either the payoff to bidder 1 is unchanged, diminished, or zero with a higher probability. This means that the overall expected payoff for bidder 1 is less.

3 Gaming Netflix

One strategy would be to first identify movies in the romantic comedy genre (this can be done by hand or according to Netflix's own classification) and look at users who reviewed such movies. We would then identify users who like these movies and model the distribution of their reviews across all movies. We could then create profiles in which reviews are drawn from these distributions to make it look as realistic as possible. Furthermore, these fake accounts would rate OMG POWNIES favorably in a manner similar to how users rate highly regarded movies, from any category.

Our method would be difficult to detect because the distribution of movie ratings would be statistically indistinguishable from how users rate any single movie. We could also take a more sophisticated approach and look at the joint distribution of ratings between movies so that we capture the positive and negative correlations between the ratings of two, three, etc. movies.

4 Google AdWords

The quality score is taking into account because it is in Google's interest (both monetarily, and to keep their users happy) to serve relevant ads that are likely to get clicked on. At a most basic level, it's clear that the quality score measures the relevancy of the ad text to the user's query. Furthermore, historical clickthrough rates can be used to reinforce ads that have been successful in the past.

Suppose that we fix everything but your quality score; that is, we fix the query, the bids of other advertisers and their quality scores. Increasing your quality score will move you up in the ad position ranking as the product (bid)*(quality score) increases. If your quality score increases but your ad position does not, the price that you pay will go down since Google charges you just enough to keep price*(your quality score) just above the bid*(quality score) of the ad ranked below you (plus one extra cent). As your quality score increases, the price required to keep your ad above the next ranked one decreases. For example, if the bid*QS of the person ranked below you is fixed, doubling your quality score nearly halves the amount that you have to pay (if it weren't for the extra cent, your price would exactly halve).

5 VCG Auction

The socially optimal pairing is to have bidder Y get slot A, bidder X get slot B, and bidder Z get slot C. Under this configuration, the value received by X and Z is 20 (18 for X and 2 for Y). If Y were not there, X would get slot A and Z would get slot B and their combined value would be 36 = 30 + 6. Thus, the harm caused by Y is 16 and that is the price they pay. The harm caused by X is 4 since with X present, Y and Z get a combined value of 52 = 50 + 2. If X doesn't participate in the auction, the combined value of Y and Z is 56 = 50 + 6. X pays 4, the harm she caused to X and Z. Finally, it's clear the Z doesn't cause any harm to X and Y as they get the same items if Z isn't present in the auction, so they pay 0.

6 The Teleporter Paradox

At equilibrium half of the cars go along the top route and half of them go along the bottom route. The latency in each case is 15. No one as an incentive to switch, as if they do, their travel time would become strictly greater than 15.

When the infinitely fast road (teleporter) is added, nothing changes. We argue that any Nash EQ has 500 cars going along AB and AC. At that point we can show no car will want to take the teleporter. Suppose there were a Nash EQ in which more than 500 cars travel along the AB link. Then more than 500 cars would also have to travel along the BD link, and they would have a travel time of more than 15, so they would have incentive to switch to the ACD path. Contradiction. Now suppose there were a Nash EQ with more than 500 cars traveling along the AC road. Any car that takes the path ACBD would rather take the path ABD as the AB time is a fixed 5 whereas the AC time is greater than 5. So there can't be a Nash EQ in which more than 500 cars take AC and some of them take the CB road. However, if all of the 500+ cars took ACD, the travel time would be greater than 15 and some of them would rather take ABD.

At this point we've argued any Nash EQ must have 500 cars taking AC and AB. If any car took the teleporter, more than 500 cars would be traveling along BD and it would be in their interest to take the CD road instead.

7 Personalized Search

We read some very good ideas. Here's some of the things that we were looking for:

Signals

- textual content of posts, tweets, etc.
- demographics info, e.g., age, location
- ratings, e.g., number of votes, likes, retweets, etc.
- time
- properties of user's social network, e.g., closeness of friends, prestige of posts
- Potential concerns: many signals depend on users' willingness to disclose information

Ranking

- simple: by time, by page rank/HITS scores, by number of votes/likes, by TF-IDF

- more sophisticated: weight various signals, e.g., recent posts have higher weight, friends' content considered more important then non-friends, etc.
- Vulnerable to: spamming, account hijacking
- ranking algorithms should run periodically, e.g., every few minutes for newest content, every few hours for all content

Aggregation

- need access to social networks (e.g., Facebook, Twitter)
- crawl them or work with them (via API)
- need streaming support for continuous arrival of information
- Privacy concerns: some users don't want their information searched

Recommendation based on

- personal interests, posts, history, etc.
- friends' interests, etc.
- important figures' interests, etc.
- collaborative filtering (e.g, similar to Netflix)
- Concern: new content may take a while to 'surface' (e.g., lack of votes)