Fall 2014 Professor John Rust
Economics 425 Georgetown University

Answers to Problem Set 4
December 11, 2014

Part I: Short Questions

1. Following up on the question raised in class, in the case of a Cournot duopoly with a linear demand function $p = a - bq$ and two firms with equal constant marginal costs of production of c and no fixed costs, is overall social surplus (i.e. the sum of consumer surplus plus total profits) higher if the firms act as ordinary Cournot duopolists or is it higher if firm 1 is a Stackelberg leader and firm 2 acts as a Stackelberg follower? Compute the equilibria in the two cases and show all calculations for full credit.

Answer: As we worked out in class, with a linear demand curve the Competitive Equilibrium (ce) quantity supplied is $q_{\text{ce}} = (a-c)/b$ (set price = mrginal cost, solving the equation $c = a + bq_{\text{ce}}$. Profits are zero in competitive equilibrium and the total social surplus equals consumer surplus, which is the area under the demand curve above the horizontal line of height c (the marginal cost of production) to the left of the competitive equilibrium quantity q_{ce}. Using your geometry that the area of a triangle is $1/2 \times$ base $\times$ height, we have base $= q_{\text{ce}}$ and height $= a - c$, where a is the intercept of the demand curve (i.e. the price at $q = 0$). Thus consumer's surplus (which equals total surplus) is $\frac{1}{2}\frac{a-c}{b}(a-c) = \frac{(a-c)^2}{2b}$.

In the *Cournot duopoly case* (cd) as I showed in class, the Nash equilibrium output of each firm is $\frac{1}{3}$ of the CE level, so that the total output, $q_{\text{cd}} = \frac{2}{3}\frac{(a-c)^2}{b}$. In the Cournot case there are positive profits, and the total profits of the two firms is $(p_{\text{cd}} - c)q_{\text{cd}} = \frac{2}{9}\frac{(a-c)^2}{b}$. Consumer's surplus is (using the formula for the area of a triangle of height $a - p_{\text{cd}} = \frac{1}{3}(a-c)$ and base q_{cd} is $CS = \frac{1}{9}\frac{(a-c)^2}{b}$. Adding total profits plus consumer's surplus to get total surplus, we get $TS = \frac{(a-c)^2}{3b}$ which is lower than in the perfectly competitive case. The difference is due to the *inefficiency* of the Cournot equilibrium, which leads to a *deadweight loss* equal to the difference between total surplus under competitive equilibrium (which is the maximum possible surplus) and total surplus under the Cournot-Nash equilibrium.

In the *Stackelberg duopoly case* (sd) we let firm 1 be the *Stackelberg leader* and firm 2 be the *Stackelberg follower.* The Stackelberg follower assumes that the Stackelberg leader can *precommit* to an output q_1^* that is best for it, i.e. the value of q_1 that maximizes firm 1's profits. The Stackelberg leader, unlike in the Cournot case, takes into account firm 2's optimal response to its choice of q_1^*, i.e. it assumes that q_2^* is not fixed at the Nash equilibrium level (as it is in the Cournot case) but rather firm 1 assumes that q_2 will be a best response to its value of q_1. In class we worked out the best response for firm 2 to a given output of firm 1: it is

$$q_2^*(q_1) = \frac{a - c - bq_1}{2b} \tag{1}$$

Thus, taking firm 2's reaction function given above into account, firm 1's output is the solution to

$$q_1^* = \underset{q_1}{argmax}\,[a - b(q_1 + q_2(q_1))]\,q_2 - cq_2 = \left[a - b\left(\frac{a-c}{2b} - \frac{q_1}{2} + q_1\right)\right] q_1 - cq_1 \tag{2}$$

Taking the derivative of this with respect to q_1 and setting it to zero and solving, we get

$$q_1^* = \frac{a-c}{2b} \tag{3}$$

Thus, the Stackelberg leader chooses $\frac{1}{2}$ of the competitive equilibrium output, which happens to be the same output a monopolist would choose. Plugging the Stackelberg leader's optimal choice for q_1^* into the Stackelberg follower's reaction function we get

$$q_2^* = q_2^*(q_1^*) = \frac{a-c-bq_1^*}{2b} = \frac{a-c}{4b} \tag{4}$$

Thus, total output in the Stackelberg duopoly case is

$$q_s = q_1^* + q_2^*(q_1^*) = \frac{3}{4}\frac{a-c}{b} \tag{5}$$

Thus, we see that total output in the Stackelberg case is $\frac{3}{4}$ of the competitive equilibrium ouptut, whereas the Cournot equilibrium total output was $\frac{2}{3}$ of the CE output. It follows that prices are lower in the Stackelberg case, and thus consumer's surplus is higher. Using the "$\frac{1}{2}\times$ base $\times$ height formula for a triangle with height $a - p_s = \frac{1}{4}a + \frac{3}{4}c - c = \frac{3}{4}(a-c)$ we get

$$CS = \frac{9}{32}\frac{(a-c)^2}{b}. \tag{6}$$

Total producer's surplus (profit) is given by

$$PS = (p_s - c)q_s = \frac{1}{4}(a-c)\frac{3}{4}\frac{a-c}{b} = \frac{3}{16}\frac{(a-c)^2}{b}. \tag{7}$$

Since $\frac{3}{16} < \frac{2}{9}$, total profits of firms 1 and 2 are less under the Stackelberg equilibrium than under the Cournot equilibrium, although the Stackelberg leader's profits are higher than the profits of each firm in the Cournot duopoly (the Stackelberg leader's profits are $\frac{1}{8}\frac{(a-c)^2}{b}$ which are greater than the profits of each of the Cournot duopolists, $\frac{1}{9}\frac{(a-c)^2}{b}$). Total surplus in the Stackelberg case is thus,

$$TS = CS + PS = \frac{15}{32}\frac{(a-c)^2}{b}. \tag{8}$$

Since $\frac{15}{32} > \frac{1}{3}$, it follows that total surplus is higher in the Stackelberg case and thus deadweight loss is lower. In fact, since $\frac{15}{32}$ is nearly equal to $\frac{1}{2} = \frac{15}{30}$, the deadweight loss in the Stackelberg case is very low. What the the Stackelberg leader has done is to take a greater share of the surplus as profits, but the competition from firm 2 causes the total output to be larger than the monopoly case, and this lowers the deadweight loss in the Stackelberg case compared to the monopoly case even though the Stackelberg leader produces the monopoly level of output. You can check that in the monopoly case the total surplus is

$$TS = \frac{3}{8}\frac{(a-c)^2}{b}, \tag{9}$$

which is lower than in the Stackelberg case, but higher than in the Cournot duopoly case.

2. Suppose Disneyland is trying to decide the profit-maximizing pricing strategy for its Disneyland theme park. Suppose each ride in the park costs c per ride in terms of electricity and labor and other costs. If Disney is a monopolist and charges a price per ride and it believes that all consumers have a

utility function for the number of theme park rides r and consumption of all other goods g of $u(r,g) = r^{.1}g^{.9}$ and the average income of a consumer is $y = 20,000$ (20 thousand dollars), then if $c = 1$ (i.e. a marginal cost of \$1 per ride), what price per ride should Disney charge if it expects that there is a market of 100,000 possible consumers who would be coming to the Disney theme park each summer? What strategy yields higher profits: a) an optimal (linear) monopoly price per ride or b) a two-part tariff consisting of a fixed entrance fee F and a constant price p to ride on the rides in the theme park? To get full credit, calculate the profits under both pricing schemes, a and b, and show which one is higher.

Answer: I screwed up this question a bit and it does not have an absolutely clear cut answer for reasons I will explain shortly. If consumers have a Cobb-Douglas utility function, they want to consume at least some Disneyland rides not matter how high the price is since if they consume zero rides, they also get zero utility. Let's ignore the constraint that one can only go on an integer-valued number of rides, and somehow allow people to consume "fractional" Disney Land rides (e.g. instead of going on a 10 minute Space Mountain ride, suppose people could choose some fraction of time to go on it for a pro-rated price, e.g. ride for 1 minute at 1/10 of the price, then in effect we can treat the number of rides r as a continuous decision.

Now I am not going to repeat the calculation, but as I showed in class and as we covered many times in the first part of the semester, with a Cobb-Douglas utility function, the demand for rides is

$$r(p,y) = \frac{.1y}{p} = \frac{2}{p} \tag{10}$$

Now, with 100,000 potential identical consumers, Disney's profits from charging a price p per ride is

$$\Pi(p) = 100000(p-c)\frac{2}{p} \tag{11}$$

It is easy to see that $\Pi'(p) > 0$, so that in effect, Disney should charge an infinite price per ride to maximize profits. To be mathematically correct, one should say that "there is no finite price that maximizes profits" and thus technically speaking there is no solution to this problem. I gave full credit to anyone who realized this. However there are several things one can do to change the problem slightly so that there is a well-defined solution to the problem, One way is to change the utility function so that if the price of rides at Disney is too high, nobody wants to go to Disneyland, which seems like a realistic assumption. So suppose instead of a Cobb-Douglas utility function, we use the following utility function

$$u(r,g) = \sqrt{r} + g \tag{12}$$

Notice now that when $r = 0$, utility is no longer zero, but is given by $u(0,g) = g$. Thus, for this utility function people do not need to go on some small fraction of a ride at Disneyland in order to be happy, no matter what the cost per ride is. If we work out the demand for rides in this case, we get

$$r(p,y) = \frac{1}{4p^2} \tag{13}$$

Actually, we see that with this utility function, the person wants to consume some small fraction of a ride regardless of how high the price is, just like the Cobb-Douglas utility function, but in this case the demand for rides declines to zero at a faster rate, $\frac{1}{p^2}$, than in the Cobb-Douglas case where demand goes to zero at rate $\frac{1}{p}$. The reason why the person want to consume some small fraction of a ride no matter how high the price is due to the fact that the utility of rides is $\sqrt{r}$ and thus the marginal utility goes to infinity as $r \downarrow 0$, so even at very high prices the consumer finds it optimal to go on some small fraction of a ride. Now with this utility function, Disney's profit maximization problem is given by

$$\max_p \Pi(p) = 100000(p-c)\frac{1}{4p^2} \tag{14}$$

Using calculus and taking the derivative $\Pi'(p)$ and setting it to zero and solving for p^*, we get $p^* = 2c$. Thus, in this case there is a finite optimal price that Disney would want to charge consumers. Now let's

consider an alternative way for Disney to charge: it charges every customer a fixed entry fee f and once in, each customer can go on as many rides they want at a price equal to marginal cost: $p = c$. Which way of pricing gives Disney higher profts? First, we can calculate profits under the optimal monopoly price calculated above:

$$\Pi_l(p^*) = 100000\frac{(2c-c)}{4(2c)^2} = \frac{100000}{16c} \tag{15}$$

Now what are Disney's profits under the fixed fee F? It gets zero profits per each ride a customer takes, so its total profits are simply

$$\Pi_{2p}(c,F) = 100000F \tag{16}$$

where Π_{2p} denotes profits under a (nonlinear) 2-part tariff and Π_l denotes profits under a (linear) optimal pricing rule. To see which of these two pricing schemes is better, we have to determine what the highest entry fee F the monopolist could charge. To figure this out, realize that the consumer always has the option of not going to Disneyland at all. If they do not go to Disneyland and spend all of their income on other goods g, their utility will be $u(0,g) = y$, since with an income of y and the price of other consumption normalized to 1 we have $u(r,g) = \sqrt{r} + g = g = y$ when $r = 0$ and $g = y$. This level of utility determines the consumers' *reservation utility:* Disney cannot charge an entry fee that is so high that people's utility after paying the entry fee is lower than the utility they can get by not going to Disneyland at all. The utility of a consumer who pays the entry fee F and can buy rides at marginal cost, $p = c$ is given by

$$V(c,y-F) = \frac{1}{2c} + y - F - \frac{1}{4c} = y - F + \frac{1}{4c} \tag{17}$$

Now setting F to make the consumer indifferent between paying the fixed entry fee F and going on rides at marginal cost $p = c$ and not going to Disneyland at all, we get

$$y = y - F + \frac{1}{4c} \tag{18}$$

so $F^* = \frac{1}{4c}$ and Disney's profits under a 2-part tariff is given by $\Pi_{2p}(c,F^*) = 100000/4p$ which is four times larger than the profits it can earn by charging a linear optimal monopoly price and no fixed fee. By allowing its customers to go on rides at a price equal to marginal cost, Disney can eliminate the deadweight loss inherent in linear monopoly price. It is able to capture and convert the deadweight loss into profits by charging the fixed fee F^* for admission. In fact, Disney has succeeded in extracting all of the surplus from the consumers and taking all of the social surplus for itself as profits.

The other way to get an answer to this problem using the Cobb-Douglas utility function is to assume that either Disney faces a competitor (e.g. Great America) that charges a price p, or there is government regulation that prevents it from charging a price higher than $p > c$. In either of these cases, Disney, if it operates with linear pricing, will want to charge the highest price it can get away with, which is p. I leave it as an exercise for you to show that with the Cobb-Douglas utility function (or indeed with any strictly convex utility function) Disney would get more profits by charging a fixed entry fee F and let consumers go on rides at marginal cost c rather than charge the price $p > c$ per ride and no entry fee. (Note: This question could show up on the final exam, so it is a good idea for you to try to work this out, at least for the Cobb-Douglas, but ideally you should be able to show this more generally for all strictly convex utility functions.)

5. Consider the following two player game. The game starts with an initial "kitty" of \$100. Player 1 can take any part of this kitty for him/herself. Whatever the player does not take gets passed on to player 2 in the next round, but the amount passed on is *doubled.* Then player 2 decides how much of the kitty, if any, is passed on for player 1 at the next round. Whatever amount is passed on in each round is doubled. The game runs for a total of 4 rounds. Thus, an example of one possible "play" of the game is for player 1 to pass the entire \$100 to player 2 in the first round. This amount is then doubled to \$200 for player 2 in the second round. If player 2 takes \$50 for him/herself at this stage and passes on \$150 to player 1 in round 3, then the \$150 is doubled to \$300 and player 1 decides how much of this to take in round 3. If player 1 takes \$200 in round 3, the remainder, \$100, is doubled, giving player 2 a total of \$200 in the 4th and final round. In this final stage player 2 could take the entire \$200 for him/herself, or give part of it to player 1. If player 2 takes all of the \$200 in this example, then player 2 gets a total of \$250 (\$50 taken in round 2 and the \$200 in round 4) and player 1 gets a total of \$200 from the \$200 he/she took in round 3. Suppose this game is played by two complete strangers who are kept in separate rooms and cannot communicate or collude in any way. If both players are rational and they don't only care about maximizing the amount they personally can earn from this game but they give some weight to how much their opponent will earn (even if the opponent is a complete stranger!), describe the Nash equilibrium outcome of this game. (Hint: the utility function for player i is $u_i(P_i, P_{-i}) = \sqrt{P_i} + \frac{1}{2}\sqrt{P_{-i}}$ where P_i is player i's monetary payoff (in total) and P_{-i} is their opponent's payoff. Use backward induction, starting from player 2's optimal decision in round 4 of the game).

Answer: We solved a "selfish" version of this problem on the practice midterm exam. Now we solve an "altruistic" version of this game, where each player gets a total utility equal to their own utility function (which is the square root of their total dollar payoff from the game) plus 1/2 of their "perceived" utility of the other player (i.e. the perceived utility of the other player is the square root of the opponent's total payoff). Note that the perceived utility of the other player is not the same as the actual payoff, since each player cares about each other. If we required that each player's total utility is their own utility plus 1/2 of the *actual* utility of their opponent, we would seem to have a much harder "fixed point" or "circularity" problem, since the opponent's actual utility includes 1/2 of the actual utility of their opponent, etc. However if we actually figure this out, the utility function is basically still the same. To see this, let $U_1(p_1, p_2)$ be the "actual" utility function of player 1, and $u_1(p_1)$ be the "subutility function" that player 1 gets from his/her own payoff, p_1. In this case we have $u_1(p_1) = \sqrt{p_1}$. Similarly let $U_2(p_1, p_2)$ be the "actual" utility function of player 2, and let $u_2(p_2)$ be player 2's "subutility" for player 2's own payoff, i.e. $u_2(p_2) = \sqrt{p_2}$. Now we have

$$
\begin{aligned}
U_1(p_1, p_2) &= u_1(p_1) + \frac{1}{2}U_2(p_1, p_2) \\
U_2(p_1, p_2) &= u_2(p_2) + \frac{1}{2}U_1(p_1, p_2)
\end{aligned}
\tag{19}
$$

This is a system of two equations in two unknowns. We can solve this to get the "reduced-form" representation of the actual utility functions as

$$
\begin{aligned}
U_1(p_1,p_2) &= \frac{4}{3}\left(u_1(p_1)+\frac{1}{2}u_2(p_2)\right) \\
U_2(p_1,p_2) &= \frac{4}{3}\left(u_2(p_2)+\frac{1}{2}u_1(p_1)\right)
\end{aligned} \tag{20}
$$

Thus, we see that the "actual" utility function is just a positive scalar multiple of the "perceived" utility function, so we will get the same results regardless of whether we analyze the problem using the actual or preceived utility functions of the two players in the game. Since the perceived utility functions are simpler (they don't have the extra $\frac{4}{3}$ factor), let's use them, i.e. we use the utility functions

$$
\begin{aligned}
U_1(p_1,p_2) &= \sqrt{p_1}+\frac{1}{2}\sqrt{p_2} \\
U_2(p_1,p_2) &= \sqrt{p_2}+\frac{1}{2}\sqrt{p_1}.
\end{aligned} \tag{21}
$$

Now, let's work out the Nash equilibrium of this 4 stage, alternating move game between the two players. I will use "P1" to identify player 1 and "P2" to identify player 2 as a shorthand. Do not confuse P1 with P1's *payoff*, which is p_1, and similarly for P2.

As usual we analyze this game by backward induction. Since the alternating moves take place quickly, we ignore any discounting and assume that the payoffs p_1 and p_2 are the *total* payoffs that P1 and P2 get from playing this game, respectively. The total payoff for P1 is the sum of 3 payments: the amount P1 takes out of the "kitty" for him/herself in round 1, x_1, the amount P1 takes out of the kitty for him/herself in round 3, x_3, and the "terminal payoff" that P2 will leave to P1 in the last round of the game. The total payoff for P2 is the sum of two payments: the amount P2 takes out of the kitty in round 2, x_2, plus the amount that P2 takes away for him/herself in the last round of the game, x_4.

Thus, to do the dynamic programming correctly, we have to set up the right "state" variable to represent the history of the game at the beginning of round 4 when P2 makes the decision about how much x_4 to take away for him/herself, and how much to leave for P1. To make this decision, P2 needs to figure out the total payoffs that P1 and P2 will get, p_1 and p_2, so that P2 can evaluate his/her utility function $U_2(p_1,p_2)$. So the required information that P2 needs is: 1) the size of the kitty V_3 that P1 passed on to P2 to divide in the last round of the game, 2) the amount x_2 that P2 took for him/herself in round 2 of the game, and 3) the amounts x_1 and x_3 that P1 took for him/herself in rounds 1 and 3 of the game. Thus the "state variable" in round 4 is (V_3,x_3,x_2,x_1) where V_3 is the size of the kitty, and (x_3,x_2,x_1) is the history of payments taken out of the kitty in the previous 3 stages of the game that have been played so far. With this information we can compute the payoffs to players 1 and 2. For P2, his/her payoff p_2 is given by

$$
p_2(x_4,V_3,x_2,x_1) = x_4+x_2. \tag{22}
$$

For P1, his/her payoff is given by

$$
p_1(x_4,V_3,x_3,x_2,x_1) = 2(V_3-x_4)+x_3+x_1. \tag{23}
$$

To see this, if the kitty available for P2 to divide in round 4 has V_3 dollars in it, then if P2 take x_4 of these dollars for him/herself, then the remaining amount, V_3-x_4, is doubled and given to P1. This is P1's "terminal payoff". But P1's *total payoff* p_1 is the sum of this terminal payoff plus the amounts P1 took out of the kitty for him/herself in rounds 1 and 3 (when P1 had the chance to take money out of the kitty).

So with the payoffs determined, P2 must solve the following problem

$$
\begin{aligned}
x_4(V_3,x_3,x_2,x_1) &= \underset{0\le x_4\le V_3}{argmax}\, u_2(p_1(x_4,V_3,x_3,x_2,x_1)) + \frac{1}{2}u_1(p_2(x_4,V_3,x_3,x_2,x_1)) \\
&= \sqrt{x_4+x_2} + \frac{1}{2}\sqrt{2(V_3-x_4)+x_3+x_1} \qquad (24)
\end{aligned}
$$

Taking the derivative of P2's payoff function on the right hand side of the equation above with respect to x_4, setting it to zero, and then solving for the optimal x_4 we get

$$
x_4(V_3,x_3,x_2,x_1) = \frac{2}{3}V_3 + x_1 + x_3 - x_2 \qquad (25)
$$

This tells us that the more that P2 took out of the kitty for him/herself back in round 2, x_2, the less P2 will take out in the final round. Why? It is because P2 is altruistic and cares about not only P2's own "subutility" $\sqrt{p_2}$ but P2 also puts $\frac{1}{2}$ weight on P1's subutility $\sqrt{p_1}$ as well. From the formula for $x_4(V_3,x_3,x_2,x_1)$ we can now deduce P1's terminal payoff: $2(V_3-x_4) = \frac{2}{3}V_3 - 2x_1 - 2x_3 + 2x_2$. We see that if P1 is "greedy" by taking out more from the kitty in stages 1 and 3 when P1 can do so, P1 will be penalized by P2 in the last stage of the game by getting a lower terminal payoff, and for each dollar that P1 tries to take out of the game for him/herself "early" P1 is penalized by a 2 dollar reduction in his/her terminal payoff by P2. This will be important to unerstand what will happen earlier in this game.

Now go back to round 3. At this stage it is P1's turn to decide how to split the kitty. The "state" of the system at round 3 is (V_2,x_2,x_1), i.e. V_2 is the amount of the kitty that P2 passed on (after being doubled, according to the rules), and x_2 is the amount of the kitty taken out by P2 in round 2, and x_1 is the amount of the kitty taken out by P1 in round 1. In round 3, P1 is deciding about x_3, the amount of the kitty to take out at round 3. If x_3 is taken out by P1, then $V_2 - x_3$ is the amount that is passed on to round 3. This amount is doubled, according to the rules, so we have

$$
V_3 = 2(V_2 - x_3) \qquad (26)
$$

So P1's decision problem is

$$
\begin{aligned}
x_3(V_2,x_2,x_1) &= \underset{0\le x_3\le V_2}{argmax}\, u_1(p_1(x_4(V_3,x_3,x_2,x_1),V_3,x_3,x_2,x_1)) + \frac{1}{2}u_2(p_2(x_4(V_3,x_3,x_2,x_1),V_3,x_3,x_2,x_1)) \\
&= \sqrt{x_3+x_1+\frac{2}{3}[2(V_2-x_3)]-2x_1-2x_3+2x_2} + \frac{1}{2}\sqrt{x_2+\frac{2}{3}[2(V_2-x_3)]+x_1+x_3+x_1-x_2} \\
&= \sqrt{\frac{4}{3}V_2 - \frac{7}{3}x_3 - x_1 + 2x_2} + \frac{1}{2}\sqrt{\frac{4}{3}V_2 - \frac{1}{3}x_3 + x_1} \qquad (27)
\end{aligned}
$$

We see that the right hand side of the last equation above is *decreasing* in x_3, so it follows immediately that $x_3(V_2,x_2,x_1) = 0$, i.e. *it is always optimal for P1 to pass on the entire kitty to P2, and trust that P2 will divide it "fairly" in round 4 of the game.* The trust that P1 has is "credible trust", i.e. P1 correctly perceives that P2 cares not only about him/herself, but also about P1's welfare, and so P2 will in fact give P1 a share of the kitty. When we also account for the fact that the kitty doubles when P1 passes it on to P2 in round 4, we see that *on the margin, each dollar that P1 passes on to P2 in round 4 will effectively be doubled and come back to P1 as a terminal payment at the end of the game.* Thus, being a rational player, P1 chooses to trust P2 and pass on the entire kitty to P2 to split in round 4.

We can continue to work backward and although the algebra is a little tedious, you can show that the optimal solutions for x_1 and x_2 are also 0. That is, in this game, the optimal solution really is to "pay it forward" since by doing so, the kitty keeps doubling and both players make out better in the end. Thus, the outcome of the game is that the initial kitty of $V_0 = 100$ is passed on to round 2 in its entirety by P1, where it doubles to $V_1 = 200$. Then P2 passes on this entire amount to round 3, where it doubles again to $V_2 = 400$. P1 then passes this amount on in its entirety in round 3, whereupon it doubles again to \$800. Since this is the last round, then P2 takes an amount $x_4(800,0,0,0) = \frac{2}{3}800 = 533.33$ and passes the remainder, \$266.66, to P1. But this amount passed on is also doubled, so that P1 actually walks away with \$533.33. Thus, both players walk away with \$533.33. Paying it forward has turned out to be very profitable indeed!

Note that this Nash equilibrium solution is *not* the one that maximizes the *total* payoff to the two players. That solution would be to pass everything on at every stage, including the last, so that P1 would have a total terminal payment of \$1600. In setting up the problem, I assumed that no "collusion" or *ex post* transfers could occur (i.e. transfers between the two players after the experiment). I have in mind here a situation where the two "subjects" in the experiment are unknown to each other and would have no way of contacting each other after the experiment was over and the payments were divided up. However if this is not the case, then another possibility is that P2 should pass the entire kitty to P1, so it becomes \$1600 and trust P1 to divide up this kitty after the experiment was over. If P1 and P2 could do this and communicate and get together and divide up the \$1600 after the experiment was over, would P2 want to do this? You can assume that P1 "owns" the money after then end of the experiment and that P2 cannot threaten P1 with harm, nor can they write any legally binding contract prior to the experiment specifying how to divide up the proceeds after the experiment. If this is the case, would it be a good idea for P1 to trust P1's altruism and pass the entire amount to P1, and then trust that P1 will give some of the \$1600 back to P2 after the experiment is over?

5. What is the output supply function for a competitive firm? Show that if the price of output i, p_i, increases, then the production of good i, y_i, cannot fall, but must stay the same or increase. (Hint: use Hotelling's Lemma and the convexity of the profit function). What can you say in general about the "output substitution effect" $\partial y_i/\partial p_j$, i.e. the effect on the production of good i of an increase in the price of output j, holding all other things equal? Can you say that this is always positive or negative? Using an example production function

$$y_1^2 + 4y_2^2 = \left[{x_1}^2 + {x_2}^2\right]^{1/2} \tag{28}$$

Compute $\partial y_1/\partial p_2$ and see what the "output substitution effect" is for this special case.

Answer: As I discussed in class, even though we know that the "own price effect" $\partial y_i/\partial p_i$ is positive for outputs and negative for inputs, we cannot unambigously say that "cross price effects" $\partial y_i/\partial p_j$ are positive or negative: it depends on the case. If outputs i and j are *complements* then $\partial y_i/\partial p_j$ could be positive (if price of good j goes up you want to make more of good j and the complementary good i). If outputs i and j are *substitutes* then $\partial y_i/\partial p_j$ could be negative (if price of good j goes up, then want to substitute more production towards increasing output of good j substituting away from subsitute output i).

For this production function we have (using reasoning that will be repeated in problem 7 below) "concave" isoquants for the inputs x_1 and x_2. This means that we will either specializing in using only input x_1 to produce y_1 and y_2, or only input x_2, depending on which one is cheaper. Let x be the cheaper of these 2 inputs and let w be its price. Then the profit maximization problem can be written as

$$\max_{y_1,y_2,x} \mathcal{L}(y_1,y_2,x,\lambda) \equiv p_1y_1 + p_2y_2 - wx + \lambda(x - y_1^2 - 4y_2^2) \tag{29}$$

where λ is the Lagrange multiplier for the production function constraint. Taking first order conditions we get

$$\begin{aligned} 0 &= \frac{\partial}{\partial y_1}\mathcal{L}(y_1,y_2,x,\lambda) = p_1 - 2\lambda y_1 \\ 0 &= \frac{\partial}{\partial y_2}\mathcal{L}(y_1,y_2,x,\lambda) = p_2 - 8\lambda y_2 \\ 0 &= \frac{\partial}{\partial x}\mathcal{L}(y_1,y_2,x,\lambda) = -w + \lambda \end{aligned} \tag{30}$$

Solving the last equation we get $\lambda^* = w$ and

$$\frac{p_1}{p_2} = \frac{y_1^*}{4y_2^*}. \tag{31}$$

Since $\lambda^* = w > 0$ the production function constraint is binding, so we can use the first two equations of the first order conditions in equation (30) above to derive the *input demand function:*

$$x^* = x(p_1, p_2, w) = y_1^*(p_1, p2, w)^2 + 4y_2^*(p_1, p_2, w)^2 = \frac{p_1^2}{4w^2} + \frac{p_2^2}{16w^2}. \tag{32}$$

Then using this input demand function combined with the first order condition for y_1^* in equation (30) above we get the following *output supply function:* for $y_1(p_1, p_2, w)$:

$$y_1^* = y_1(p_1, p_2, w) = \frac{\sqrt{x(p_1, p_2, w)}}{\sqrt{1 + \left(\frac{p_2}{4p_1}\right)^2}} = \frac{\sqrt{\frac{p_1^2}{4w^2} + \frac{p_2^2}{16w^2}}}{\sqrt{1 + \left(\frac{p_2}{4p_1}\right)^2}} \tag{33}$$

Using this formula, we can now compute $\partial y_1^*/\partial p_2$ and see if we can determine if it is positive or negative (i.e. if outputs y_1 and y_2 are complements or substitutes). By doing some algebral (sorry, I am not going to type this here going through the calculus in all detail, but as an outline of how I did this calculation note that we can write

$$y_1^* = \frac{N}{D} \tag{34}$$

where N denotes the numerator term in equation (33) and D denotes the denominator term. But numerator and denominator terms are functions of p_1, p_2 and w, of course, but letting N' denote the partial derivative of the numerator term with respect to p_2 and letting D' similarly denote the derivative of the denominator with respect to p_2, then by calculus we have

$$\frac{\partial y_1^*}{\partial p_2} = \frac{N'}{D} - \frac{ND'}{D^2} \tag{35}$$

and using the rule for derivative of square root function and the chain rule of calculus, we have

$$\begin{aligned} D' = \frac{\partial D}{\partial p_2} &= \frac{p_2}{16p_1^2}\frac{1}{D} \\ N' = \frac{\partial N}{\partial p_2} &= \frac{p_2}{16w^2}\frac{1}{N} \end{aligned} \tag{36}$$

Substituting these equations into the equation for $\partial y_1^*/\partial p_2$ given in equation (35) above and doing some algebraic simplification, we get

$$\begin{aligned}\frac{\partial y_1^*}{\partial p_2} &= \frac{p_2}{16w^2}\frac{1}{DN} - \frac{p_2}{16p_1^2}\frac{N}{D^3} \\ &= \frac{N}{D}\left[\frac{p_2}{16w^2}\frac{1}{N^2} - \frac{p_2}{16p_1^2}\frac{1}{D^2}\right] \\ &= y_1^* p_2\left[\frac{1}{4p_1^2+p_2^2} - \frac{1}{16p_1^2+p_2^2}\right] > 0. \end{aligned} \tag{37}$$

From the final expression for $\partial y_1^*/\partial p_2$ we can see this derivative is positive (i.e. y_1 and y_2 are complementary outputs) since the denominator of the first term in the last bracketed expression in equation (37) is smaller than the denominator in the second term in the brackets.

6. Consider a firm selling mufflers. Each day there is a probability p that exactly 1 customer will come to the store to buy a muffler. Suppose the retail price of the muffler (the price the firm can sell to the customer) is p_r and the wholesale price of a muffler (the price the firm can buy mufflers from the manufacturer at) is p_w. Naturally we assume that $p_r > p_w$ so the firm makes profits from selling mufflers. Suppose that each time the firm orders more mufflers to replenish its inventory, it incurs a fixed transport cost K *regardless of how many mufflers it buys from the manufacturer.* Suppose there is also a storage/holding cost of mufflers and if the firm has q mufflers in its inventory, it costs c per muffler to store them. Suppose the firm is an infinite-horizon profit maximizer and its discount factor is $\beta \in (0,1)$.

a. What is the profit maximizing inventory strategy for this firm? Write down the Bellman equation for the firm's optimization problem and characterize the nature of the solution for full credit.

answer Let q be the (integer-valued) inventory of mufflers. We write the Bellman equation for $V(q)$ the optimal value (expected present discounted value of profits) for the muffler company. For $q > 0$ we have

$$V(q) = \max\left[V_n(q), V_o(q)\right], \tag{38}$$

wherer $V_n(q)$ is the value of not ordering more mufflers and $V_o(q)$ is the value of ordering more mufflers, given by

$$\begin{aligned} V_n(q) &= \left[p * p_r - cq + \beta(pV(q-1) + (1-p)V(q))\right] \\ V_o(q) &= \max_{q'>0}\left[p * p_r - cq - K - p_w q' + \beta\left(pV(q+q'-1) + (1-p)V(q+q')\right)\right]. \end{aligned} \tag{39}$$

and for $q = 0$ we have

$$V(0) = \max\left[\beta V(0), \max_{q'>0}\left[-K - p_w q' + \beta\left(pV(q'-1) + (1-p)V(q')\right)\right]\right]. \tag{40}$$

There is a slightly different version of the Bellman equation when $q = 0$ since when $q = 0$ the firm is "stocked out" so it cannot sell anything if a customer arrives. Thus, in the Bellman equation for $V(0)$ there is no term $p * p_r$ for the expected sales revenue, since the firm has no inventory and thus it is unable to sell a muffler to any customer who arrives. Further, the inventory holding cost cq is also zero in this case.

Besides writing the Bellman equation, I looked to see if you provided an equation characterizing the optimal ordering strategy (optimal decision rule for ordering inventory) $d(q)$. As I discussed in class, under certain conditions the optimal ordering strategy can be of the "(S,s)" form, that is, there are two integers $S \geq s \geq$ such that

$$d(q) = \begin{cases} S - q & \text{if } q \leq s \\ 0 & \text{otherwise} \end{cases} \tag{41}$$

What this means is that if inventory on hand falls below the lower threshold s then it is optimal to order a quantity $d(q) = S - q$ that is sufficient to return total inventory $q + d(q)$ back to the "target level" S.

If fact, as you might have discovered from solving this problem numerically, the optimal decision rule is *not* of the simple (S,s) form in this case. In this problem I have assumed (via the way I wrote the Bellman equation) that when the firm orders mufflers, there is a *one day delivery lag*, so any new mufflers the firm orders, q', will arrive the *next day*.

Alternatively, I could have assumed that there is *no delivery lag*, i.e. when the firm orders new mufflers, the new mufflers *arrive the same day, in the morning before any customer arrives.* Under this *instantaneous delivery assumption* there is a slightly modified Bellman equation given below. The basic equation (38) still holds, but the equations for $V_o(q)$ and $V_n(q)$ need to be slightly modified as follows

$$\begin{aligned} V_n(q) &= [p * p_r - cq + \beta[pV(q-1) + (1-p)V(q)], \\ V_o(q) &= \max_{q'>0} \left[p * p_r - c(q+q') - K - p_w q' + \beta\left(pV(q+q'-1) + (1-p)V(q+q')\right)\right]. \end{aligned} \tag{42}$$

The equations above cover all values of q¿ 0.

$$V(0) = \max\left[\beta V(0), \max_{q'>0}\left[p * p_r - K - p_w q' - cq' + \beta\left(pV(q'-1) + (1-p)V(q')\right)\right]\right]. \tag{43}$$

This equation differs from equation (40) (in the case where there is a 1 day delivery lag) because the firm can guarantee there is never any unserved customer (no stockouts) when it has instantaneous delivery of new mufflers. Thus, even if $q = 0$ at the start of the business day, with immediate delivery, the firm can order a muffler and have it available to sell by the time the shop opens, to any customer who might show up.

Since I did not specify in the problem whether you should assume immediate delivery or delivery with a one day lag, I accepted either formulation of the Bellman equation given above. When there is a delivery lag, the optimal decision rule is no longer of the (S,s) form. Instead it is of the "(S_0,S,s) form", i.e. the decision rule is given by

$$d(q) = \begin{cases} S_0 & \text{if } q = 0 \\ S - q & \text{if } 0 < q \leq s \\ 0 & \text{otherwise} \end{cases} \tag{44}$$

Thus, the firm sets a potentially different optimal inventory target, S_0, if it is ordering "from scratch" (i.e. with $q = 0$) compared to when it is *re-ordering* when it has a postive inventory level already, $q > 0$. In the latter case the target inventory level is S rather than S_0. I did not expect you to *prove mathematically* that this is the case, but rather to see from the numerical solutions that this is what form the solutions take. From the numerical solutions, it is possible to develop *general mathematical proofs* that the optimal decision rule is of the (S,s) form when there is immediate delivery, and of the (S_0,S,s) form when there is a one day delivery lag.

b. Suppose that $\beta = .99$, $c = 0.1$, $p_w = 100$, $p_r = 150$, and $K = 100$. Calculate the optimal policy for the firm numerically, using a computer and report the present discounted value of the firms profits assuming it has in inventory $q = 5$ mufflers.

answer I omitted a key parameter here, p, the probability that a customer arrives. Clearly the level of inventory that should be ordered will also depend on this parameter. It should be clear that the orders will be the highest when $p = 1$, since then the firm will periodically "stock up" to meet demand for a period of days before reordering. The firm is balancing the fixed cost of placing an order, $K = 100$ against the inventory holding cost per muffler, c. With immediate delivery, it should be clear that $s = 0$, the firm waits until it is stocked out before ordering more mufflers. When mufflers arrrive with a one day lag, the firm will order when it has one muffler left, $s = 1$, since otherwise it would risk a "stock out" and not have any mufflers to sell when a customer arrived. When $p = 1$ the optimal policy is to order $S = 13$ mufflers when $q = 0 = s$ when there is immediate delivery. When there is a delivery lag, the optimal policy is to order $S_0 = 13$ mufflers when the firm is stocked out, and $S = 13$ mufflers when it has $q = 1 = s$ mufflers in inventry. This means that if no customer arrives on the day it orders when it has $q = 1$ muffler left in stock, then it will have a total of $14 = S + 1$ mufflers on the next day, whereas if it had $q = 0$ mufflers, it would order $S_0 = 13$ mufflers and have 13 mufflers at the start of the next business day. So there is a slight difference in the optimal strategy depending on whether there is immediate delivery or not.

I wrote the Matlab programs `setup.m` `bellman.m` and `succapp.m` to calculate the optimal policy. `setup.m` just sets up the parameters and global variables needed to solve the problem. `bellman.m` is a Matlab function that calculates the Bellman operator, i.e. it evaluates equation (38) above given any input value V, i.e. it evaluates $\Gamma(V)$ where Γ is the "Bellman operator" and the Bellman equation amounts to a fixed point of the Bellman operator

$$V = \Gamma(V) \tag{45}$$

In this case the "state space" is the level of inventory q and it takes integer values $q = 0, 1, 2, \ldots$ so we can treat V as a vector in some finite dimensional Euclidean space R^n since we can guess that the firm will not want to keep more than a finite number n of mufflers in stock at any time. We will guess that $n = 20$ is an upper bound on the number of mufflers that the firm would ever want to order or have on hand (if this is too small and we find that $S \geq n$ then we can increase n and resolve until we have a value of $n > S$, so we are confident that our arbitrary guess about what n is does not impinge on the solution of the problem).

Now recall the method of *successive approximations* to solve for the unique fixed point $V = \Gamma(V)$. We make an initial guess $V_0 = 0$, i.e. a vector of zeros in R^n. Then we stick this intial guess into the Bellman operator to get

$$V_1 = \Gamma(V_0) \tag{46}$$

which is an updated guess of the value function. We keep doing this, resulting in a sequences of value functions $\{V_t\}$ given by

$$V_{t+1} = \Gamma(V_t) \tag{47}$$

and we stop these iterations when $\|V_{t+1} - V_t\| < \varepsilon$ for some small convergence tolerance ε. The program `succapp.m` is a Matlab program that carries out this successive approximations algorithm using the Bellman operator as programmed by `bellman.m`. I set a convergence tolerance $\varepsilon = 0.000001$. At this small tolerance I have "almost" solved the Bellman equation and from the

Bellman equation I can uncover the optimal decision rule $d(q)$. The Matlab program notation for the value function is `V` and the optimal decision rule is `dr`. By inspecting these you can determine the optimal inventory ordering rule. You can download and run these programs to get the numerical answers below.

Figure 1 plots the value function, the approximate solution to the Bellman equation (38) and approximate fixed point $V = \Gamma(V)$ for the case of instantaneous delivery of orders when the arrival probability is $p = 1$. In this case, the policy is of the (S,s) form with $s = 0$ and $S = 13$.

Figure 1: Value function for inventory problem with instantaneous delivery when $p = 1$

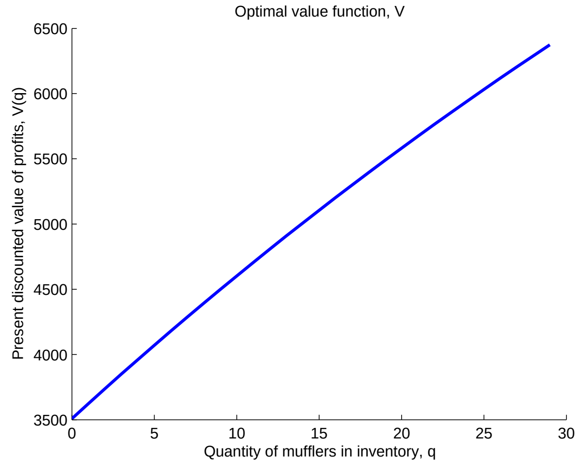

In particular, the expected discounted value of future profits for the company when $q = 5$ is $V(5) = 3208.95$.

Figure 2 plots the optimal order quantity S as a function of the customer arrival rate p and it is an increasing function of p as expected. We see that when $p = 1$ then $S = 13$. The smallest arrival probability p where it makes sense for the firm to re-order mufflers is $p = .1309$. For any arrival probability less than this, it is optimal for the firm to shut down and sell off any mufflers it has in stock but never reorder any more. As p rises above this threshold the smallest value of S is $S = 5$ until the arrival probability reaches $p = .175$ when S increases to $S = 6$, and so forth in a series of steps until it reaches $S = 13$ when $p > .902$.

Figure 2: Optimal order quantity, S, as a function of arrival probability p

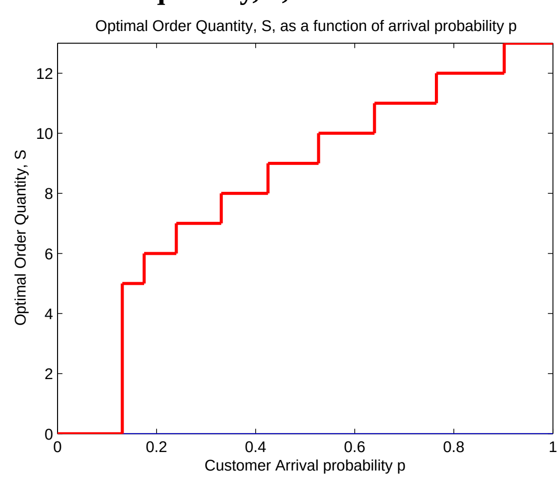

c. Suppose there are occasional opportunities to buy mufflers at a lower wholesale price than $p_w = 100$. Suppose with probability $q \in (0,1)$ the firm can buy as many mufflers as it wants at a price of $p_{wl} = 75$. Write the Bellman equation in this case. What are the state variables for this problem? What are the decision (control) variables? How does the nature of the solution change compared with the solution you characterized in part a above?

answer Let q denote the quantity of mufflers in inventory as before, but now add an additional state variable x to indicate whether the firm can buy mufflers at the lower price of $75 per muffler (compared to the normal price of $100 per muffler). Thus $x = 0$ denotes the case where there is no sale on mufflers, so the company faces the high wholesale price $p_w = 100$, and $x = 1$ denotes the case where there is a sale on mufflers so the company can buy them at the lower wholesale price of $p_w = 75$. Then the value function is $V(q,x)$ and is given by

$$V(q,x) = \max\left[V_n(q,x), V_o(q,x)\right], \tag{48}$$

where, as above, $V_n(q,x)$ is the value of not ordering, and $V_o(q,x)$ is the value of placing the (optimally sized) order for more mufflers. These values are given by

$$\begin{aligned} V_n(q,x) &= p * p_r - cq + \\ &\quad \beta\left[q\left(pV(q-1,1) + (1-p)V(q,1)\right) + (1-q)\left(pV(q-1,0) + (1-p)V(q,0)\right)\right] \\ V_o(q,x) &= \max_{q'>0}\left[p * p_r - cq - K - p_w(x)q' + \beta q\left[pV(q+q'-1) + (1-p)V(q+q')\right]\right. \\ &\quad \left. + \beta(1-q)\left[pV(q+q'-1) + (1-p)V(q+q')\right]\right]. \end{aligned} \tag{49}$$

In this Bellman equation $p_w(x)$ denotes the (state-dependent) wholesale price of mufflers, so $p_w(0) = 100$ and $p_w(1) = 75$. The Bellman equation above is for the case of a one day delivery lag. With instantaneous delivery the Bellman equation can be modified similar to our discussion in the answer to part a above, and the Bellman equations for $V(0,x)$ can be derived similarly.

In the case of instantaneous delivery, the optimal ordering strategy can be shown to take the form of a *generalized* (S,s) *rule*. In a generalized (S,s) rule the values S and s are functions of the state variable x, and so can be written as $(S(x), s(x))$. This result can be proven using the argument in Hall and Rust (2007) *Economic Theory*. Howevr I did not expect you to have seen (or even guessed) that there might be a generalization of the (S,s) inventory strategy that might be applicable here. I mainly wanted to see that you could write down the correct Bellman equation.

d. Calculate the optimal strategy for the firm numerically using a computer for the modified version of the problem in part c, assuming $q = 0.05$.

answer I wrote Matlab programs `bellman2.m` and `succapp2.m` to solve this problem using the successive approximations strategy described in the answer to part b above. The function `bellman2.m` implements the Bellman operator in this two-dimensional problem. The value function V can be stored as a *matrix* of dimension $n \times 2$ where the two columns of V represent the value in the case $x = 0$ and $x = 1$ respectively. Similarly the optimal decision rule $d(q,x)$ is stored in the $n \times 2$ matrix `dr2` which is returned implicitly by `bellman2.m` as a global variable.

Figure 3 below plots the optimal decision rule $d(q,s)$ for the case of instantaneous delivery with $p = 0.35$ (i.e. there is an 35% probability that customer arrives to buy a muffler every day). We see that when $s = 0$ (the wholesale price of mufflers is high, $p_w = 100$) that it is never optimal to order mufflers. Instead the firm only orders mufflers when they are on sale, i.e. when $x = 1$. In that case we have $S(1) = 14$ and $s(1) = 6$. Thus, when mufflers are on sale, the firm will not order unless $q < s(1) = 6$, and when it does order, it orders enough to reach the target inventory level $S(1) = 14$. Note that $S(0) = s(0) = 0$, so that the firm does not order new mufflers whenever the wholesale price is high, i.e. when $p_w(0) = 100$.

Figure 3: Optimal order quantity, S, as a function of arrival probability p

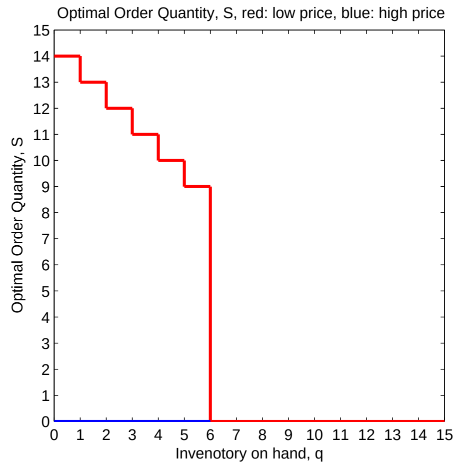

However when we increase the arrival probability, say to $p = 0.8$ we get a different optimal generalized (S,s) strategy. Now we have $S(1) = 24$ and $s(1) = 11$ and $S(0) = 8$ and $s(0) = 0$. This implies the possibility that it is optimal for the firm to occasionally stock out of mufflers, even if it can get them delivered instantly. For example suppose the firm initially has a stock of 24 mufflers purchased at the low muffler price of $p_w(1) = 75$. Suppose the price returns back to the high price, so that the state becomes $x = 0$ for a long stretch of time. When $x = 1$ it is not optimal for the firm to reorder mufflers until $q = 0 = s(0)$ and then it orders only $S(0) = 8$ mufflers. That is the firm experiences a stock out and only orders relatively few mufflers to tide it over (and avoid losing opportunities to sell mufflers to customers) while it is waiting for another reduction in the wholesale price of mufflers when it would be optimal for it to “buy big” and purchase enough mufflers to reach the target level $S(1) = 24$. Then as the firm starts to sell off this inventory of mufflers purchased at the low wholesale price, it will reorder again if $q < 11$ and there happens to be a sale going on (i.e. if $x = 1$). Otherwise it will continue to sell off its inventory of mufflers until all of them are sold and then it will only order the smaller quantity $S(0) = 8$ to tide it over until another low price opportunity comes along.

7. Consider a seller trying to sell a painting at an auction. Suppose that there are N buyers who will participate in an auction if the seller holds one. Suppose that (normalized to millions) that the seller knows that the valuations of buyers are random draws from a distribution on the $[0,1]$ interval (where 1 now denotes $1 million dollars, and the cumulative distribution function of the values is

$$F(v) = Pr\{\tilde{v} \leq v\} = v^2 \tag{50}$$

a. What is the expected amount a single buyer would be willing to pay for this painting?

answer This question can be answered following the discussion in my lecture notes on auctions and I will not take the trouble to provide a complete answer here.

b. Write a formula for the bidding function that the bidders would use in a symmetric Bayesian Nash equilibrium of a *first price auction* for the painting. What is the probability that the painting will be sold, and what is the expected price that the seller will receive for the painting if there are $N = 5$ bidders participating in the auction?

c. Suppose instead that the seller runs a *second price auction.* What is the value that the seller can expect to receive for the painting when there are $N = 5$ bidders participating in this auction?

d. Suppose there are $N = 3$ bidders with valuation for the painting equal to $v_1 = .2$, $v_2 = .8$ and $v_3 = .5$. What price will the seller receive if a) she adopts a first price auction for the painting, b) she adopts a second price auction for the painting, or c) she adopts an *all pay* auction?

e. Suppose that the seller adopts a second price auction, but sets a *reservation price* for the painting of $r = .2$. That is, the seller will not sell the painting unless the highest bid is at least $r = .2$ ($200,000). Describe how the use of the reservation price affects the buyers' bidding strategies in this auction, if at all? What is the probability that the seller will sell the painting when there is a population of $N = 5$ potential bidders whose true values for the painting are given by a probability distribution with cumulative distribution function $F(v) = v^2$? Can you calculate the expected revenue the seller will receive? If so, which is better for the seller, to have no reservation price, or to set a reservation price of $r = 2$?