2. (Coherence and Dutch book *[20 total points]*) On 2 Apr 2001 a senior writer for the web page *Sportsline.com*, Mark Soltau, posted an article about the Masters golf tournament that was about to be held on 5–8 Apr 2001. Among other things he identified the 24 players (among the 93 golfers in the field) who were, in his view, most likely to win the tournament, and he posted odds *against* each of them winning (for example, his quoting of 10–1 odds on Phil Mickelson meant that his personal probability that Mickelson would win was $\frac{1}{1+10} \doteq 0.091$), which are summarized in Table ?? below.

(a) If the 24 odds quoted by Mr. Soltau were taken literally, show that the personal probability specification implied by his posted odds was incoherent. (In fact Mr. Soltau may well have been quoting un-normalized odds, which is a fairly common practice in sports, but let's take him literally in this part of the problem.) *[5 points]*

(b) It would be nice to demonstrate Mr. Soltau's incoherence by explicitly providing a set of bets that would be guaranteed to lose him money, but that's actually fairly complicated (hint for the previous part of this question: that's *not* what I had in mind for you to do in (a)). To take a simpler example that has the same flavor as Mr. Soltau's mistake (if his odds are taken literally), pretend that he's handicapping (setting odds for) a tournament in which only Tiger Woods, Phil Mickelson, and some other unnamed golfers are playing, and he announces 3 to 1 odds in *favor* of Woods winning and 1 to 1 odds in favor of Mickelson (again without specifying any odds for the other golfers). (To be clear on the relationship between odds and money, here's how it works in horse-racing (and Mr. Soltau would have to play by the same rules): suppose that a bookie at the horse track offers odds of 4 to 1 against horse A, and I bet (say) \$1 on that horse to win; if horse A wins I enjoy a net gain of \$4, otherwise I suffer a net loss of \$1.) Work out an explicit set of bets to offer Mr. Soltau that would constitute a Dutch book against him. If Mr. Soltau were willing to accept arbitrarily large bets, is there any theoretical limit to the amount of money you would be guaranteed to win from him? Explain briefly. *[10 points]*

(c) In practice sports bookies only allow people to make bets *for* individual golfers, so that in reality you're not allowed to construct a wager like {\$$x$ on Woods to win and \$$y$ on Mickelson to lose}. Can you make Dutch book against Mr. Soltau under these conditions? Explain briefly. *[5 points]*

3. (Bayes' Theorem; based on problem 7 in chapter 1 of Gelman et al. *[20 total points]*) In the old television game show *Let's Make a Deal*, there are three doors; behind one of the doors is a car, and behind the other two are goats, with the assignment of prizes to doors made at random. You — the contestant, who prefers cars to goats — are asked to pick a door. After you choose (of course you can do no better than picking at random), the emcee, Monte Hall, who knows where the car is, opens one of the other doors to reveal a goat, and he offers you the choice of staying with the door you originally picked or switching to the other unopened door. Suppose that Monte Hall uses the following algorithm to decide which door to reveal to you after you've chosen (say) door 1. If the car is behind door 2 he shows you door 3; if it's behind door 3 he shows you door 2; and if it's behind door 1 he randomizes between showing you doors 2 and 3 with equal probability. Should you switch or stay with your original choice?

(a) Explicitly use Bayes' Theorem to work out the chance of winning the car under each strategy. *[10 points]*

(b) How would you explain intuitively to someone who favors the inferior strategy why the other one is better? *[10 points]*

4. (Conditional probability, and review of the normal distribution; based on problem 4 in chapter 1 of Gelman et al. *[20 total points]*) (American) football (not soccer) experts provide a *point spread* (PS) for every football game as a measure of the difference in ability between the two teams. For example, team A might be a 3.5–point favorite over team B. This means that the proposition that A (the favorite) defeats B (the underdog) by 4 or more points is considered a fair bet, i.e., $P(A \text{ wins by more than 3.5 points}) = \frac{1}{2}$. If the PS is an integer, the implication is that A is as likely to win by more points than the PS as it is to win by fewer points than the PS (or to lose); there is a positive probability that A will win by exactly the PS, in which case neither side is paid off. In Chapter 1 Gelman et al. present data on the PS and actual game outcome for 672 professional football games played during the 1981 and 1983–84 seasons, and they show that the histogram of the quantity (actual outcome – PS) is well approximated by a normal distribution with mean 0.07 and standard deviation (SD) 13.86, suggesting that a good predictive distribution for the actual result of an NFL football game would be normal with mean equal to the PS and SD 14 points (two touchdowns). (If you're in the habit of betting on NFL games this should give you pause, e.g., if a team is favored by a touchdown the chance it will win, according to this uncertainty assessment, is only about 69%.) It turns out that there were 12 games in this data base with PS values of 8 points, and the actual outcomes in those games were −7, −5, −3, −3, 1, 6, 7, 13, 15, 16, 20, and 21, with positive (negative) values indicating wins by the favorite (underdog). Consider the following conditional probabilities:

$$P(\text{favorite wins}|\text{PS} = 8)$$
$$P(\text{favorite wins by at least } 8|\text{PS} = 8)$$
$$P(\text{favorite wins by at least } 8|\text{PS} = 8 \text{ and favorite wins})$$

(a) Estimate each of these using the relative frequencies of the games with an 8–point PS. *[5 points]*

(b) Estimate each using the normal approximation to the distribution of (actual outcome – PS). (You can use a normal table from any statistics book, or the error function `erf` in `Maple`.) *[10 points]*

(c) Which of these approaches to uncertainty assessment seems to have produced better answers here? How should we define "better"? Explain briefly. *[5 points]*

5. (Cromwell's Rule and its implications for Bayesian learning *[10 total points]*) Prove the following two facts: for any D such that $P(D) > 0$,

(a) If $P(A) = 0$ then $P(A|D) = 0$.

(b) If $P(A) = 1$ then $P(A|D) = 1$.

In the usual application of these facts (as in the HIV case study in class), A is a proposition whose truth value is unknown to You (such as the HIV status of the patient) and D represents some data relevant to A (such as the result of a screening test like *ELISA*); in this setting (a) and (b) together are referred to as *Cromwell's Rule* (I'll give the history behind this in class). What are the implications of Cromwell's Rule for the use of Bayes' Theorem as a formal model for learning from data? Explain briefly.