6. (8 points) Beatriz used logistic regression on a data set derived from people living in Framingham, MA to learn a linear logistic classifier $\sigma(\theta^T x + \theta_0)$ giving the probability that an adult with features x will develop heart disease in the next decade.

   Her friend, John, would like to use the same logistic regression classifier (i.e., the θ^* and θ_0^* learned by Beatriz) to make predictions for people living in Norway. However, he notices that heart disease is much less common in Norway and thinks that the model may need to be adjusted to account for this.

   (a) Consider a specific patient with feature vector x. How could John adjust θ_0, relative to the θ_0^* learned by Beatriz, so as to make smaller the probability of this patient developing heart disease?

   Solution: Assuming all other parameters remain the same, John would need to make θ_0 smaller, i.e. $\theta_0 < \theta_0^*$.

   (b) John realizes that choosing the right value of θ_0 is tricky since he doesn't have access to any labeled data from Norway. John tells Beatriz that he only plans to use the model to find the 10% of individuals with highest probability of developing heart disease so that he can closely follow them and make sure they are tested appropriately.

   "Aha!", says Beatriz. "In that case, any value of θ_0 would suffice, and you can simply make use of my original linear logistic classifier!" Explain why Beatriz is right.

   Solution: Since σ is a monotonic function and θ_0 is a constant that does not depend on x, the ranking of patients according to $\sigma(\theta^T x + \theta_0)$ is the same no matter what the value of θ_0 is.

   Side remark: one could also multiply θ and θ_0 by any constant strictly greater than 0 and we would still get the same ranking.